

Linking Hierarchical Conjoint with Hierarchical Bayes Regression

by Peter Kurz

Abstract

Over the past few years we have successfully used hierarchical conjoint designs in a number of commercial applications. The use of Hierarchical Conjoint in applied settings and as a subject of research-on-research has shown, that they are often a superior alternative to partial profile and traditional discrete choice designs especially when a large number of attributes is being investigated. The paper presents how to set up hierarchical conjoint and how to estimate the utilities for the combined model.

1. Why to use Hierarchical Conjoint

Over the last two decades, conjoint analysis has become widely accepted for research involving product development, as well as measuring the sensitivity to pricing changes. The objective of conjoint analysis is the measurement of individual level part-worth utilities for product features or components, and then the estimation of preference for “defined” products through market simulations. The accuracy of these conjoint simulations and usability for market share estimates has been a topic of debate almost as long as conjoint has been founded (Wittink and Walsh, 1988).

Full profile conjoint- or choice-based trade-off studies have traditionally been limited to no more than six attributes. Partial profile, self-explicated scaling or hybrid methods have been used when a large number of attributes need to be included in the model. Full profile studies allow for the estimation of interaction terms and generally present more realistic choices to the respondent than partial profile or self-explicated approaches. However, clients often want to test a long list of potential product features that may or may not be included in the final product, depending on the results from research. These product features are often cognitive in nature. Additionally, they may be interested in complex pricing issues that require some interaction effect estimation or wish to test certain attributes such as brand and price in a “full profile” format. Being limited to six attributes renders traditional full profile trade-off analysis useless in these situations.

Therefore in every days live, many conjoint applications must deal with lots of attributes, usually more than can be presented effectively in full profiles at one time, both full profile conjoint or discrete choice analysis are limited to a small number of attributes. But however clients often want to include a long list of product features. One form of conjoint analysis that has received much attention, and has been the topic of many papers during the last years, is Sawtooth Software's Adaptive Conjoint Analysis (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. But 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.

Another technique for many attributes paid much attention in literature during the last years was Partial Profile Choice Based Conjoint (see Patterson & Chrzan 2004). Respondents may have the same difficulties keeping in mind that all other attributes not involved in the current question are assumed to be equal. Again the same effect than in ACA can be shown, the importance particularly for high-involvement categories is underestimated due to spreading respondents' attention across individual attributes. One effect can be an underestimate of the importance of price (especially if many attributes included) and occasionally, even with the most efficient fractional factorial design, we still end up with more products than can be practically accommodated.

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. The result that can occur is the effect of attribute additivity and affects any trade-off method with a large number of attributes in. This effect is also known as the

“Number of Attributes Effect” For Example the more important attribute “engine type” – when to choose a car – has less influence on the purchase decision than the addition of two minor attributes like “illuminated ashtray” and a “clipboard for notes”.

This Number of Attributes Effect influence both ACA and Partial Profile Choice Based Conjoint. This would tend to lessen the importance of the more important attributes, of which price is probably one. The solution for both type of problems shown above is the hierarchical conjoint approach.

The approach to solving this problem, and providing a more accurate determination of the importance of less important attributes, has been hierarchical conjoint analysis. In hierarchical conjoint, each respondent participates in two or more conjoint studies. Each set of attributes is treated like its own trade-off study. A fractional factorial design is created for each set of attributes. Respondents are asked to rate or rank two or more smaller sets of attributes rather than one large set. The utilities are calculated for each trade-off exercise independently and linked together to create one final set of utilities. One conjoint study is used to examine the more important product attributes, and a second, third and so on, conjoint study with the same respondents concentrates on the less important attributes. The first design is used to quantify all of the relevant attributes and features other than the one which are on lower level and attributes which represents the micro ones with simple attribute levels; the second measures the less important attribute in more details. Typically, the second study is used only to solve for the attributes with lower hierarchy in the decision process, relative to other attributes. Because the second study is used to solve for fewer parameters, less information is required per respondent.

There are several ways to implement hierarchical 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, discrete choice experiments, are preferred for the first study. 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 ratings-based or hybrid conjoint method (like ACA) as the second method, or simply a self-explicated questioning. Alternatively, one could use a choice-based conjoint or discrete choice method on the second level too.

2. Idea of the Hierarchical Approach

Jordan Louviere (1992) assumes that in complex choice situations respondents tend to split up the choice tasks into major and minor attributes. Under the assumption, that the preference generation of the respondents takes place in a hierarchical manner and in the same direction (macro – micro), he assumes a functional dependency between the different levels of choice.

According to this assumption the hierarchical approach splits up the attributes in a macro and one or more micro levels. The macro level attributes are the ones where respondents preference generation takes place and have a higher influence on the overall purchase decision. The micro level ones are taken only into account after the macro level ones generate the purchase decision. They are quasi nice to have features but not really relevant for creating the shortlist of choice. The basic idea behind this kind of technique is to describe the various attributes of the so called micro conjoint with the levels of one or more attributes on the macro level.

Critical to the selection of the appropriate trade-off techniques is the issue of which type of attributes, cognitive or non-cognitive, are being represented in the trade-off exercise. Cognitive attributes are attributes that are based on rational, conscious, generally verbal decision making. Such easily quantifiable factors as speed, horse power, interest rate or weight are typically cognitive. Non-cognitive attributes are attributes that are less explicit, more emotional or even less conscious such as brand, price, design, equipment or even graphic design components, such as colour, typeface or geometric shapes. One might argue that the selection of a life bank account, a laptop, a car or a water heater are all cognitive decisions and that the selection of a beer, wine, whisky, a hair spray or a pair of pants are all non-cognitive. One might also argue that all decisions made by humans are at least partially non-cognitive. This issue, as it relates to trade-off analysis, deserves further research.

However, trade-off techniques that employ direct questions (self-explicated and most hybrid techniques) all necessarily assume that the attributes being modelled are all cognitive, because at least some of the product features are being rated in a way that requires both awareness and honesty from the respondent. That is, the respondent must be aware of the degree to which a product feature affects his or her purchase decision and also be willing to admit to that degree of effect.

Additionally, any data collection methods that rely on verbal or written descriptions of product features all assume that the attributes being modelled are cognitive, because the process of understanding a verbal or written description is itself a cognitive behaviour. Non-cognitive trade-off models, that is, models involving non-cognitive attributes, should be based on an indirect trade-off technique (conjoint or discrete choice) and data collection that relies on complete product experience rather than language describing single attributes alone to communicate the product choices. That means, for trade-off models involving non-cognitive attributes, use full-profile techniques. For example, if you are modelling the car buying process, show respondents a variety of cars that they can see and touch. A consumer may respond to the phrase "a luxury car" very differently than he or she would to a particular car with a premium equipment. So keep in mind that it's important for the hierarchical approach to have the right trade-off techniques on the micro and macro level.

3. The Hierarchical Approach

A method has been developed and successfully applied numerous times which offers several advantages over both traditional full profile and partial profile conjoint and choice methods:

- A large number of product features (40 or more) can be included in the model
- Selected first order interactions can be estimated at both the micro and the macro levels
- Since product combinations are specified, via traditional experimental design, before the interview takes place, physical exhibits can be easily incorporated into the interview

The basic steps of the procedure are as follows:

- Questionnaire content:
 - Macro-Level: Discrete Choice Experiment, respondents are asked to choose one product of a list of products (choice task). Thereby the attributes and levels of the products are varied by experimental designs.
 - Discrete Choice Experiment, Adaptive Conjoint, Full Profile Conjoint or a self-explicated approach in dependence of the attributes are chosen on the micro level.
- Analysis
 - Using any of a variety of available conjoint or choice software, utility weights for each feature in the trade-off exercise can be estimated for each respondent (based on regression techniques).

- Utilities are then linked from micro-level with macro-level. On a per respondent basis, a scalar can be estimated using the common features in micro and macro level. Numerous algorithms for linking the both parts together exist. We typically use hierarchical bayes regression techniques to calculate the scalar.
- The scalar reduces the feature scores in micro level to a scale equivalent with macro level utility weights.
- On a per respondent basis, this scalar is multiplied by each score in micro level to achieve utility weights comparable to macro-level data utility weights.
- Micro-level and macro-level utility values are then merged to create one set of linked utility values (with the utility values from macro level used for the attributes spread for both levels).
- These merged utility values define the conjoint or choice model from which all subsequent simulations will be based.

If you use an self-explicated approach all attributes are necessarily additive. This can yield to misleading results if too many features are included in the model. A correction for this excessive feature bias has been developed based on the assumption that, when selecting products, respondents consider no more than six features at a time.

The correction procedure is as follows. When calculating total product utility for each individual:

- Identify the six attributes with largest relative importance
- Include only the selected levels from those six attributes when calculating total product utility for that individual
- Include all individuals' total product utilities in the simulator as normally done

Or experience shows, that such corrections are only necessary for self-explicated approaches on the micro-level, the other techniques discrete choice experiments and adaptive conjoint analysis don't need such corrections.

4. An Example formulating Attributes and Levels

Over the past few years we have successfully used hierarchical conjoint designs in a number of commercial applications. The use of Hierarchical Conjoint in applied settings and as a subject of research-on-research has shown, that they are often a superior alternative to partial profile and traditional discrete choice designs especially when a large number of attributes is being investigated. The hierarchical approach splits up the attributes in a macro and one or more micro levels.

The basic idea behind this kind of technique is to describe the various attributes of the so called micro conjoint with the levels of one or more attributes on the macro level.

For instance use engine type as attribute in the macro conjoint and number of cylinders, horsepower and fuel efficiency as attributes in the micro conjoint. The "engine type" attribute-levels than look like: 8 cylinders, 400 hp, 10 litres per 100 km as best and 4 cylinders, 100hp 18 litres per 100km as worst and one ore more levels in between (graphic 1). The procedure is as follow: estimate the utilities on the micro and macro level separate and than scale the micro level in such a way, that the combined levels of the micro conjoint are equal to the respective level of the macro attribute (graphic 2).

Macro level:

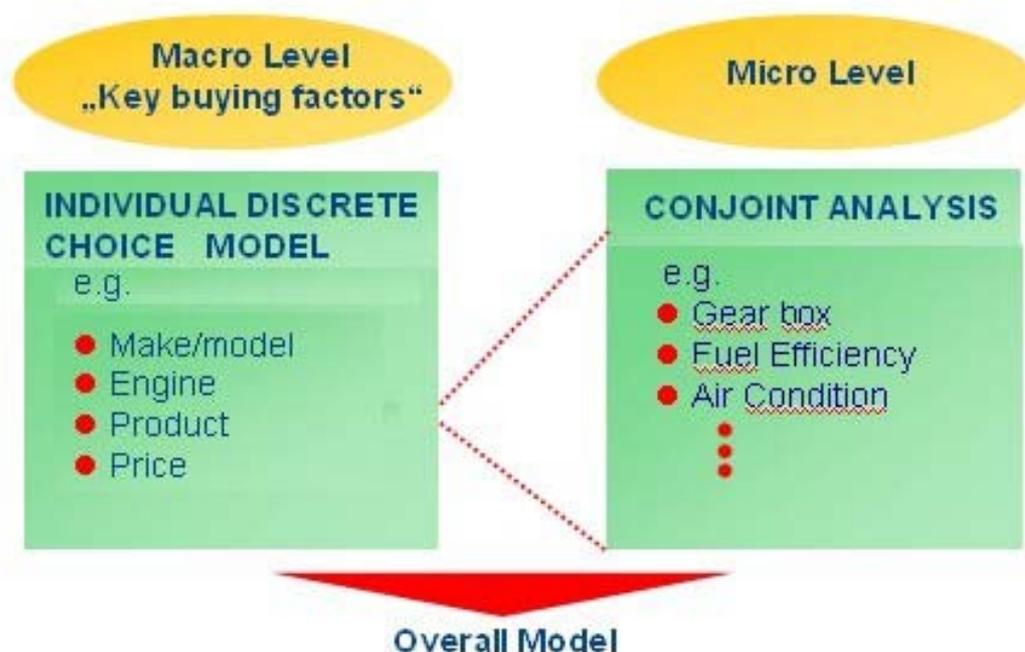
- Make/model:
 - Mercedes
 - BMW
 - Audi
 - Volvo
- Engine:
 - 4 cyl/150 hp
 - 6 cyl/180 hp
 - 6 cyl/200 hp
- Product:
 - Gearbox: manually
 - Fuel efficiency: 12 L/100 km
 - No Air condition
 - Gearbox: automatic
 - Fuel efficiency: 8 L/100
 - With Air condition
- Price:
 - \$ 24.000
 - \$ 26.000
 - \$ 28.000

Micro level (Product):

- Gearbox:
 - Manually
 - Automatic
- Fuel efficiency:
 - 8 L / 100 km
 - 10 L / 100 km
 - 12 L / 100 km
- Air condition:
 - No
 - Yes

Graphic 1: Formulating Levels for hierarchical conjoint studies

Another alternative to describe the attributes of the micro conjoint on the macro level is to use only phrases for the macro attribute levels, for example basic equipment package, medium equipment package and luxury equipment package. In this cases you have to take care that the respondents have a clear expression what the level descriptions mean. Preferably you than have to show the micro-level conjoint first to train them what the various attributes on micro level mean and than explain the new macro levels on the bases of the micro attributes. The macro conjoint should be done at the final stage of the procedure. One problem is, that if you have to much micro conjoints in your study, the main phase where respondents preference generation takes place when respondents are already tiered and are overburdened thru the whole questioning procedure.



Graphic 2: Combining Micro- and Macro-Level

5. Linking the Micro- and Macro-Conjoint

Various techniques of linking the micro and macro level together are described in literature during the last years. Starting with very simple techniques like using the spread and calculate a scaling factor for each level and end up with more advanced ones like various regression approaches. A more advanced idea is to use hierarchical bayes regression techniques to fit the two parts together.

Just like in choice based conjoint we can equate a hierarchical bayes regression equation, which equates all of the attribute levels to the response given. But than instead of solving for the features we can substitute the utilities for everything without the combined attribute, leaving only the combined attribute levels and the respective scaling factor to be estimated. Than we minimize the scaling factors with the iterative procedure of the hierarchical bayes technique on the individual (respondents) level.

Alternatively we can solve for a single scaling factor for the micro level utilities.

A comparison between the simple techniques – linking the utilities with one single scaling factor and the bayesian approach shows clearly the advantages and disadvantages of each approach and give a first impression which techniques too use. The datasets used for the comparisons are simulated data on the one and real studies with different degree of complexity and out of different markets on the other hand. Methods will be initially compared by their ability to predict individual choices against holdout tasks. Each technique produces a different set of utility values that may be used to predict respondents choice and compare it against holdout tasks. We access the success of each technique by the average number of successful predictions, as well as the pairwise comparison of both the macro and the macro level predictions.

6. Simulated Studies

TNS Infratest has recently completed a large methodological study to compare alternative implementations of hierarchical conjoint. The study compared DCM on macro level with two methods, both of which use ACA as the second level, and which used either ratings-based conjoint or choice based conjoint at the second level.

The datasets used for the comparisons are on the one hand simulated data and on the other hand real studies with different degree of complexity and taken out of different markets.

The five different simulated data sets represents: one macro and two micro conjoints in each simulation. The macro conjoint was constructed with 7 Attributes with 1x7;2x5;1x4;3x3 levels, the micro conjoints both have the same structure and consists of 7 attributes with 5 levels each. Two of the three level attributes of the macro conjoint represents the linking attributes which can be described as good, medium and bad level. The order, structure and spread of these attributes is used to test the linking procedures.

The differences in the simulated data sets are:

1. Perfect consistent Choices in Micro & Macro Level: all objects are simulated in that way, that the answers of the micro and the macro conjoints are identical.
2. Noise in either Micro or Macro Level: we add a uniform random distribution as noise to the answers of 10% of the micro or 10% of the macro objects, to get an impression how the estimation of the utilities and the linking procedures are influenced through imperfect data.
3. Noise in both Micro & Macro Level same procedure but noise both micro and macro level.
4. Reversals in either Micro or Macro Level: In this simulation we turned about 10% of the respondents answers in the opposite direction between the micro and the macro level to built up inconsistency between the micro and macro answers.
5. Reversals in both Micro & Macro Level more complex inconsistencies by having reversals on both the micro and macro level of the linking attributes - which is a pretty good impression of real data sets.

Hit rates of the simulation are shown in table 1.

Method	Spread	Regression	HB Scaling	HB Utilities
Hit Rate % n=800				
Simulated Data without error	95,8	91,3	98,4	99,2
10% Noise in Micro-Conjoint only	91,3	73,9	93,2	94,1
10% Noise in Macro-Conjoint only	94,4	71,2	92,7	93,5
10%Noise in both Micro & Macro Level	89,7	69,8	91,8	92,9
Reversals in Micro Level	90,2	61,2	91,1	91,7
Reversals in Macro Level				
Reversals in Micro & Macro Level	69,1	58,9	86,4	87,8

Table 1: Hit rates of the simulation studies by linking technique

We clearly see, that HB Utilities - which means, that we have dropped the macro level utilities and substituted them thru the micro level ones by creating a new choice file for the combined conjoint and estimate the utility values for the whole number of attributes at one time - outperforms all the other techniques. Second best technique is the method which calculates a scaling factor to combine the three conjoints with the hierarchical bayes technique. The real surprise was the fact, that the simple calculation of a scaling factor based on the spread of the macro attribute was nearly as good as the much more complicated techniques and outperform the regression technique. The only problems occur when you have reversals on the macro level what means, that the simple coefficient fails and urgently need one of the more advanced approaches or a preprocessing of the reversals. The hierarchical bayes methods are superior techniques when you have noisy data or reversals in - what is more common in the real datasets.

7. Commercial Studies

Over the past few years we have successfully used hierarchical conjoint designs in a large number of commercial applications. The focus of this paper lies on the comparison of 4 selected studies.

1. Study from the airline sector conducted face-to-face/capi on German airports (n=1.600). Consists of 1 macro conjoint and 3 micro conjoint with an overall number of 30 attributes with up to 7 levels. All conjoint are based on discrete choice experiments.
2. Study conducted in Travel Agencies face-to-face/capi (n=600). 1 macro and 1 micro conjoint both discrete choice experiments with 9 attributes macro and 5 attributes micro with up to 9 levels.
3. Study from the Finance Sector, conducted online 600 owners of current accounts from a large savings bank in the United Kingdom. Macro is done as discrete choice experiment 9 attributes with 3 to 11 levels and 2 micro choice based conjoint with 5 attributes with 5 levels each.
4. Study face-to-face/capi 400 owner of a postpaid cell phone contract. 1 Macro as discrete choice experiment with 7 attributes and up to 8 levels; micro conjoint as adaptive conjoint with 14 attributes and 5 levels each.

Frequently, hierarchical conjoint studies include holdout tasks only at the upper level. For this test designs, we included holdout tasks in both the upper and lower level conjoint tasks. This design allows the evaluation of discrete choice models alone versus a hierarchical design. Since the utilities for discrete choice experiments were estimated using a main effects design only, 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 second level attributes like lounge quality, animation packages, statements of account, games, mp3-files.

The different 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 (hit-rates). The choice tasks used for this purpose were interviews with repeated choice tasks once at the beginning and once at the end so that we could measure the reliability of respondent choices themselves and then compare the estimated choices against the holdout tasks.
- Second, calculating the correlations between the macro and the micro attributes after linking the parts together.
- Finally, aggregate shares of preference were compared using a preference simulation model.

The three different Methods are used to measure the validity of the macro level as stand alone conjoint and again to measure the differences after linking the micro level attributes in the overall model.

Method Hit Rate %	Spread	Regression	HB Scaling	HB Utilities	Macro only	Macro first vs. last task
Airline Study n=1600	65,4	56,9	77,7	83,4	89,3	89,1
Travel Agencies n=600	68,7	61,2	83,2	84,1	86,2	87,2
Finance n=800	62,3	67,8	82,8	83,9	89,2	85,3
IT-Telecoms n=400	78,1	69,3	81,4	82,8	89,8	90,7

Table 2: Hit rates of the different studies by linking technique

Again the hit-rates show clearly, that hierarchical bayes technique outperforms the others clearly. The main difference between the simulations and the real data is, that the simple technique based on spread does not do a good job with this datasets. We conclude that noise and reversals like they occur in real datasets have more influence on the simpler methods than on the hierarchical bayes based methods - which could be the advantage of the iterative learning of the algorithms. We also can see, that the hierarchical bayes method with substitution of the micro attributes and calculating utility values for the overall model comes closest to the macro hit-rates and the hit-rates of the two control tasks in the study. Another finding is, that sometimes the consistency of the two tasks isn't as good as the consistency of the overall experiment. That show, that more research is needed on the validity of holdout tasks. To us it seems with the knowledge of our previous studies, that not always the single answer on a hold-out-task is better than the up to 20 answers of a discrete choice experiment. Another finding is, that the micro holdout tasks were predicted even better than the macro ones. This shows, very clearly that the more simple tasks are predicted much better than the complex ones from the macro level. The hood-outs are also predicted with higher hit-rates after the linking procedure. This fact can be used as an additional indicator that the linking procedures do a good job. Again the hierarchical bayes approaches outperform the others.

Method Correlations	Spread	Regression	HB Scaling	HB Utilities	Macro first vs. last task
Airlines	0,773	0,521	0,721	0,732	0,84
Travel Agencies	0,790	0,646	0,846	0,853	0,89
Finance	0,712	0,621	0,775	0,789	0,81
IT-Telecoms	0,838	0,734	0,859	0,899	0,95

Table 3: Correlations between Micro and Macro-Conjoint by linking techniques

Looking at the correlations between the different methods show nearly the same picture than the hit-rates but is much closer together between the different methods.

Again the traditional regression-method does a bad job and again we can see, the differences between the hierarchical bayes estimation results and the correlation between the two control tasks is marginal.

In most of our studies hit-rates are not the best way to measure the goodness of fit of the models, our clients normally are more interested in the prediction of shares. For measuring success in prediction of shares we use mean absolute error (MAE)

between the holdout tasks and predicted shares of choice. The share predictions were done using the share of preference model. Of course that's only half of the truth, we have no "real world" market shares for the products and can only compare the differences of the four techniques based on holdout tasks. Future research should be done in addition on datasets where we also have market share data for comparison from real market.

The hierarchical bayes linking techniques were again superior to conventional regression techniques and the method using the spread. However, among those different combinations there was an interesting reversal. The method with the estimation of a scaling factor with hierarchical bayes had the most successful (smallest) MAE and outperforms the method with substituting the attributes and a hierarchical bayes run on the overall model (see table 4). We have seen the same result with other data sets: the reason seems to be the reversals between the micro and the macro conjoint. Our simulated data shows this fact more clearly. The new estimation of the combined model can handle those reversals better on individual level and help to improve the hit rates. A simple scaling factor don't do a good job in handling this reversals - but seems to scale the utilities in a way that better predicts shares. We think, that it's very similar to the findings, that using constraints usually helps hit rate, but sometimes degrades share predictions.

Method	Spread	Regression	HB Scaling	HB Utilities	Macro first vs. last task
MAE					
Airlines	6,21	8,31	4,99	5,32	3,24
Travel Agencies	6,34	8,17	4,92	5,27	3,67
Finance	6,38	7,94	4,68	5,16	4,34
IT-Telecoms	5,89	7,89	4,34	4,95	3,16

Table 4: Mean absolute errors (MAE) by linking technique

8. Findings

The use of Hierarchical Conjoint in applied settings and as a subject of research-on-research has shown, that they are often a superior alternative to partial profile and traditional discrete choice designs especially when a large number of attributes and/or price is being investigated. HB on a combined set of micro and macro utilities outperforms all the other techniques. HB on a single scaling factor only is nearly as good as the more complicated technique and is the less time consuming approach. Scaling on base of the range of variation is a fairly good and easy alternative, if the number of reversals and the proportion of noise in micro & macro conjoint is low. But it fails if inconsistencies are present in the data. The traditional multinominal logit regression technique, do the worst job and is very vulnerable if heterogeneity is found in the data. It strikes me that there are a number of key points to take away from these analyses. It is not evident from this limited number of example if these problems we found with the regression technique and the spread calculation are specific to these methods, or if all linking methods can be improved through adjustments.

It is not clear whether holdout tasks are an sufficient criteria to find out which methods are best or if share simulations are the better values. Holdout tasks are especially in hierarchical and partial-profile conjoints a problem, cause if we decide not to show the whole set of attribute as full-profile, because it's to complex and burdens the respondent to much, same problem appears with holdout tasks. If we use holdout tasks from the macro and micro conjoints separately we don't learn the whole true. On the other hand it needs some more research to find out if hold-outs are a really good predictor or if the precision of 14 to 20 answers during the conjoint are not more reliable?

A next step in research at TNS should be the comparison of hierachical conjoints from past years with real data gather from other sources like panels and real market figures and see if the findings would hold.

The Author

Peter Kurz is Senior Consultant Data Analysis, at TNS Infratest Munich, Germany

Email: peter.kurz@tns-infratest.com

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.

François Pierre, MacLachlan Douglas L., Jacobs Anja (1991), "Bridging Designs for Conjoint Analysis: The Issue of Attribute Importance", Leuven, Belgium, Unpublished paper.

Green, Paul and V. Srinivasan (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.

Louviere, Jordan J. and Timmermans, H. J. P (1992), "Testing the external validity of hierarchical conjoint analysis models of recreational destination choice", *Leisure Science* 14(3) p. 179–194

McCullough, Dick (2000), "Comment on Kilroy and Williams" in *Sawtooth Software Conference Proceedings*, Ketchum, ID: Sawtooth Software.

Patterson Michael and Chrzan, Keith (2003), "Partial Profile Discrete Choice: What's the Optimal Number of Attributes", In *Sawtooth Software Conference Proceedings*, Sequim, WA

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.