

**PROCEEDINGS OF THE
SAWTOOTH SOFTWARE
CONFERENCE**

October 2010

Copyright 2011

All rights reserved. No part of this volume may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without permission in writing from
Sawtooth Software, Inc.

ADDED VALUE THROUGH COVARIATES IN HB MODELING?

PETER KURZ

TNS INFRATEST FORSCHUNG GMBH

STEFAN BINNER

BMS MARKETING RESEARCH + STRATEGY

SUMMARY

We re-analyzed ten CBC data sets, comparing the use of covariates in HB to standard HB runs that assume single multivariate-normal populations. With HB using covariates, respondents are not shrunk toward one common distribution. Instead, part-worths of respondents with different characteristics have different multivariate-normal distributions. From a theoretical point of view this would seem more appropriate. We tried to find out whether in practice the use of covariates offers gains in predictive validity with respect to holdout choices and real market data.

We found few significant gains in predictive validity when including covariates - no matter whether the covariates were demographics, cluster segments, or segments based on past behavior and purchase intention questions.

As a second question we analyzed whether covariates could stabilize the estimates when there are reduced numbers of respondents and choice tasks. But reducing the amount of data for HB estimation in either way did not affect outcomes much. Also when analyzing within different segments we didn't find meaningful differences in outcome with and without covariates.

Our last section deals with whether covariates could improve matters when using proportional sampling within small segments of the population. We found that covariates couldn't resolve previously identified problems with proportional sample structure. In the small segments, there were small improvements with covariates, but results were far inferior to those of proportional sampling.

INTRODUCTION TO COVARIATES

The hierarchical Bayes (HB) model is called hierarchical because it models respondent preferences as functions of an upper-level (averaged across sample) model and a lower level (individual respondent-level) model. At the lower individual-level, the respondent is assumed to choose concepts by maximizing the sums of part-worth utilities as specified in the multinomial logit model.

In the standard HB approach, the upper level model assumes that respondents are drawn from a single multivariate normal distribution, with part-worths (β_i) distributed with means α and covariance matrix D , $\beta_i \sim \text{Normal}(\alpha, D)$, where i indicates the single respondent. In HB applications, the upper-level model plays the role of a prior distribution when estimating each respondent's part-worths, and the lower-level model provides the likelihood for the estimation. Because it leverages information from the upper-level population parameters α and D , HB is able to estimate relatively stable part-worths for each individual, even when the data set provides only relatively sparse information.

The assumption of a single multivariate normal population is troublesome to some researchers, who consider most markets to be composed of distinct segments. Many researchers have considered ways to modify the upper-level model so as to be more compatible with the assumption of discrete segments. We review some of those attempts before looking at covariates.

In standard HB approaches a simple assumption is used: respondents are drawn from a single population of normally distributed part-worths. While this assumption may seem to be very simple from a theoretical point of view, it performs well in most of our studies. The single-normal-population assumption is often an influencing factor only at the start-up of the estimation and does not affect the final part-worth estimates to a large extent. Especially, it does not constrain the final part-worths to be normally distributed. HB results represent a combination of the upper- and lower-level models for each individual. If enough information is available at the individual level, the resulting part-worth utilities don't show much influence of the upper-level model compared to the impact of the lower-level model. Unfortunately the influence of the upper-level model results in some Bayesian shrinkage toward the population mean value, which tends to smooth the distribution, with a tendency toward the normal distribution (especially if information is sparse for an individual). But again, if a substantial number of choice tasks are available in relation to the number of parameters to be estimated, the Bayesian shrinkage is usually small and therefore doesn't affect the result very much.

In some instances, practitioners see problems with the assumption of a single normal population:

- The assumption that respondents are drawn from a single normal population seems for many practitioners and clients unrealistic; they assume more complicated functions.
- In segmentation studies, practitioners have expressed concern that distances between segment means are shrunk because HB tends to shrink individual estimates towards the population mean.
- In situations in which a segment of respondents (with different preference structures) is oversampled, this Bayesian shrinkage can bias the estimates for the segment means as well as the overall population means – especially for the proportional part of the sample (Fuchs, 2007).

In recent years, researchers have proposed ways of solving this problem. One is the idea from Sentis and Li (2001) of estimating for segments separately, to avoid the shrinkage problem. The idea is based on the fact that, if we use a different population mean value for each segment, shrinkage to the overall population mean would no longer occur. Sentis and Li studied seven actual CBC data sets, systematically excluding some of the tasks to serve as holdouts for internal validation. They estimated the utilities in four ways: first by using the entire sample within the same HB estimation routine (one population mean value); second by segmenting respondents according to industry sectors and estimating HB utilities within each segment (mean value for each segment); third segmenting respondents using a K-means clustering procedure based on first stage HB utilities, and then re-estimating within each segment using HB; and fourth by segmenting respondents using Latent Class and then estimating HB utilities within each segment. They found that none of their attempts to improve results by fitting subgroups separately improved predictions of holdouts.

Howell, (2007) proposed respondent weighting in HB as a solution to the disproportional sampling problems outlined in the Fuchs (2007) paper. He investigated the severity of the problem, and used simulated data to demonstrate that when subgroups are dramatically oversampled, it causes the means of smaller groups to shrink disproportionately toward the larger groups. This could bias the sample means for the proportional (under-represented) groups, and violates the accuracy of preference share simulations. Howell shows that much of the problem is due to diverging scale factors between smaller and larger subgroups. The scale for the oversampled (disproportional) groups is expanded, leading to stronger pull on the overall sample mean. The article shows with artificial data that normalizing the scale post hoc can largely control this issue. Also it concludes that implementing a simple weighting algorithm within HB (computing a weighted alpha vector) can potentially improve matters further when there are extreme differences in sample sizes between subgroups. Practitioners often deal with the problem by using disproportional sampling and estimating the groups separately for each of the small segments to avoid Bayesian shrinkage.

Other approaches to solve the problem employ multiple upper-level distributions. It is apparent that continuous heterogeneity (normal mixture) alone is superior to discrete heterogeneity (Latent Class), up through a fairly large number of segments (Rossi, Allenby & McCulloch 2005); also, a correlated (random) coefficients specification for the normal mixture is superior to an uncorrelated one; and more than one segment can be used in the normal mixture model. But for multiple mixtures of upper-level models, more complex mathematical functions and estimation procedures are needed as well as a lot of prior knowledge about the data structure to gain the right multivariate-normals. Allenby & McCulloch (2005) found that extending HB to accommodate multiple distributions leads to only minimal gains in predictive accuracy. From a practitioners' point of view we can solve this problems by estimating the relevant segments separately if we know the data structure upfront.

Advanced HB practitioners have recommended that in many cases “well-chosen” covariates could provide additional information and therefore improve parameter estimates and preference share predictions (Lenk, et.al. 1996). Covariates could be seen as another term using additional independent variables that may affect part-worths. Often, we think of covariates such as demographics like gender, age, income, company size, geographic location, etc. Unfortunately, these variables often have only low correlations with the preference structure of our choice context. The most useful covariates bring exogenous information (additional information which is not already available in the choice tasks) to the model to improve the estimates of part-worth and improve preference share predictions.

More formally, instead of assuming respondents are drawn from one normal distribution with mean vector α and covariance matrix D , an HB model which uses covariates in the upper-level model assumes that respondent part-worths are related to the covariates through a multiple regression model:

$$\beta_i = \Theta' z_i + \varepsilon_i \text{ where } \varepsilon_i \sim \text{Normal}(0,D)$$

where Θ is a q by b matrix of regression parameters, z_i is a q vector of covariates, and ε_i is a b vector of random error terms. The part-worths are now drawn from a normal distribution with mean values $\Theta' z_i$, different for each respondent. No longer shrinking the individual estimates to a single population mean α , this method shrinks them to the conditional mean Θz_i given the

subject’s covariates. With this solution the multiple regression upper-level model can use observed, segment basis variables (e.g. Country, Car Segment, Distribution Channel, etc.) to improve the estimation of the part-worths and may increase the distinction between segments in the data set.

In the standard HB model with the single normal population mean value, there are $b + [b(b+1)]/2$ parameters to be estimated in the upper-level model, where b is the number of part-worths for each individual respondent. When including covariates in the upper-level model, there are $bq + [b(b+1)]/2$ parameters to be estimated in the upper model, where b is again the number of part-worths and q is the number of parameters introduced by the covariates. If Country was the covariate, consisting of China, Russia, Italy, UK, US and Germany, q would equal 6. Including covariates in the upper-level model doesn’t alter the number of parameters estimated for covariance matrix D . Using covariates is more parsimonious than separating the sample by country and running standard HB with a different covariance matrix for each of the six separate samples. In our data set number 6, the vector of sample means plus the covariance matrix require $49 + [49(49+1)]/2 = 1274$ parameters to be estimated for the upper model in each sample, for a total of $6 * 1274 = 7644$ parameters if samples are estimated separately. Estimating as a single run with a dummy-coded covariate for country requires only 1519 parameters in the upper-level model, a very substantial saving over the number required when running standard HB within the separate segment samples. So, using covariates requires estimating many more parameters than with the single-normal “standard” model, but many less than when estimation is done separately for subgroups. Using covariates also takes more computer time than the standard single normal approach, but less than making separate runs for each segment. For more information about the mathematics of the HB model with covariates introduced, please see Orme & Howell (2009).

DESCRIPTION OF ANALYZED STUDIES

For the purpose of this paper 10 commercial studies with a total of approximately 30,000 interviews covering nearly all industries and topics were analyzed. These studies cover B2B as well as B2C markets and were conducted in almost all parts of the world. All of these studies were carefully designed and used Sawtooth Software. As to good research practice and to ensure valid results, the sample structures of these studies were disproportional, ensuring sufficient sample size for all segment cells.

	Industry	Target Group	N =	Interview Delivery	# Coice Tasks	# Est. Parameter	Concepts/ Task	Conjoint Method	Type of Natural Segments
Study 1	Tyres	B2C	189	CAWI	16	19	4	STD CBC	4 customer types
Study 2	DIY	B2C	500	CLT/CAPI	12	8	5	CBC ASD	5 distribution channels
Study 3	Adhesive	B2C/B2B	888	CAWI	8	40	16	CBC ASD	2 customer types
Study 4	Fuel Cells	B2C	926	CAWI	16	69	4	CBC ASD	3 countries
Study 5	Adhesive	B2C / B2B	600	CAPI	16	13	5	CBC ASD	5 customer types
Study 6	Automotive	B2C	8900	CAPI	15	49	5	STD CBC	6 countries x 12 car segments
Study 7	Automotive	B2C	8400	CAPI	14	86	3	CBC ASD	12 car segments
Study 8	Automotive	B2C	9200	CAPI	19	21	6	CBC ASD	6 countries
Study 9	Technology	B2C	3000	CAWI	15	74	8	CBC ASD	4 countries x 3 customer types
Study 10	FMCG	B2B	750	CAPI	18	84	12	CBC ASD	5 distribution channels

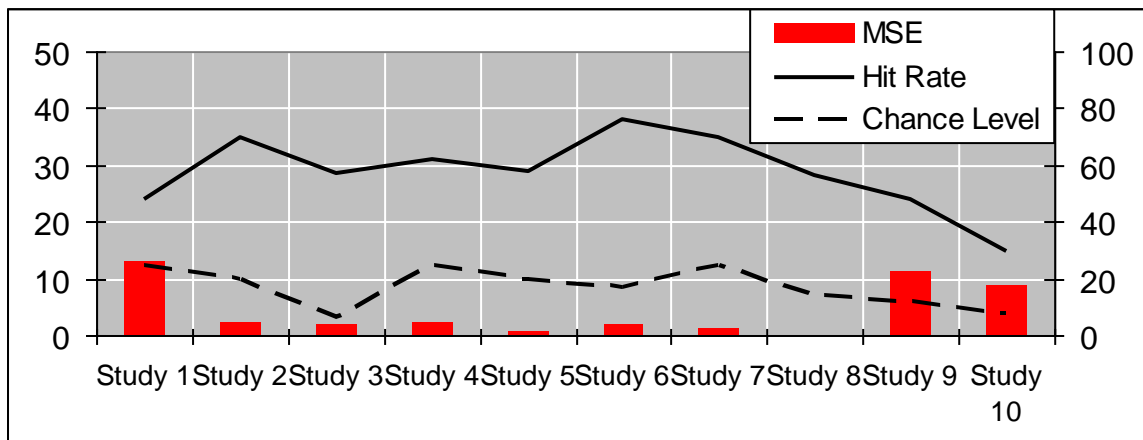
Table 1: Overview of studies analyzed

The 10 studies have been selected as they represent different scenarios from practical modeling work with conjoint analysis.

Study 1	Tires	Study on motorcycle tires in one country. Due to the different bike types in the study there is quite a lot of heterogeneity among the respondents
Study 2	DIY	Price Conjoint which was conducted in four different distribution channels which have different competitive environments
Study 3	Adhesive	This study analyzed the impact of branding on price elasticity in both, professional and private end user markets
Study 4	Fuel Cells	Study about energy supply in RVs. Very complex model with different product alternatives and 69 parameters to estimate. Difficult to recruit target group led to comparable small sample size in each of the three countries analyzed
Study 5	Adhesive	Brand/price conjoint in market with highly fragmented customer segments (some B2B, some B2C)
Study 6	Automotive	Automotive Study with different car features 6 countries 12 car segments from small mini cooper up to a lager limousine and suv's – sportscars face2face computer assisted 15 choice tasks 49 parameters Number of respondents compared to number of parameters should result in very stable estimates.
Study 7	Automotive	Automotive newer concepts of engines hybrid, active hybrid gas engines 12 segments one country more parameters asd model
Study 8	Automotive	Tires 9000 6 countries small number of parameters less concepts
Study 9	Technology	Technology 4 countries 3 typs of customers flatscreens 3000 online 74 parameter
Study 10	FMCG	10 fmcg b2b manufacturer of chips respondents retailers capi 84 parameters 5 channels hyper markets to groceries

Table 2: Description of studies analyzed

We examined both the hit rate for predicting the holdout choice, as well as the mean square error (MSE) of the base case simulation against the holdout task results. When looking at this performance measure standard HB (in the following labeled as HB STD) showed rather satisfying results in regard to these two measures:



Graphic 1: MSE and Hit of standard HB in ten studies

TYPES OF COVARIATES USED FOR ANALYSIS

For the systematic analysis of the ten commercial studies, different types of Covariates were defined:

Type 1: Membership in Natural Segment

Demographic or product specific segments (categorical variables) were used as covariates. These were for example countries, customer groups, product segments or distribution channels. For further analysis in this paper we labeled models with this type of covariate as HB COV-N (HB with covariate based of natural segment), models with independent HB estimations for every single natural segment were labeled as INDV HB.

Type 2: Membership in Segments

For this group of covariates we used Latent Class or benefit segments (categorical or dummy coded). The benefit segments were derived by cluster analysis of individual utilities. Covariates based on Latent Class segments are called HB COV-L, those based on benefit segments HB-COV-U.

Type 3: Added Data

For a limited number of studies additional data was available and used as Covariates. Such added data included purchase intention (e.g. stated budget for new vehicle) or past behavior (e.g. purchase price of last vehicle)

All ten studies were analyzed with the these types of estimation models

HB STD	Standard HB for the whole study sample
HB COV-N	HB with covariate (defined by natural segment)
INDV HB	Independent HB Estimations per natural segment
LC	Latent Class (Sawtooth Software)
HB COV-L	HB with covariate (defined by LC segments)
HB COV-U	HB with covariate (defined by STD HB utility cluster)


All estimations were performed with Sawtooth CBC/HB (v5.2.2) using standard settings (20,000 iterations, prior variance 2, degrees of freedom 5).

OBSERVATIONS DURING THE ESTIMATIONS


We could observe during the estimations that the models showed in first 1,000 - 10,000 draws a different behavior in convergence, while in the end the models converged to the same parameters than without covariates. We assume that this is caused by an influence of the upper-levels model when using covariates.

Following convergence plots demonstrate the slightly different behavior (shapes) at the beginning while finally converging towards the same parameters.

Without Covariates:




With Covariates:




RESULTS OF DIFFERENT ESTIMATIONS

Only in two of ten studies HB COV showed significantly better MSE results than STD HB. The reason for this might be the relatively large number of parameters and small samples for those two studies (study nine and ten). However in these two studies HB COV was not better than INDV HB. Looking at the Hit Rates we observed the same results.



Graphic 2: MSE and Hit Rates of different simulation models

On the other hand there was no real champion among the alternative estimation methods based on segment membership or Latent Class:



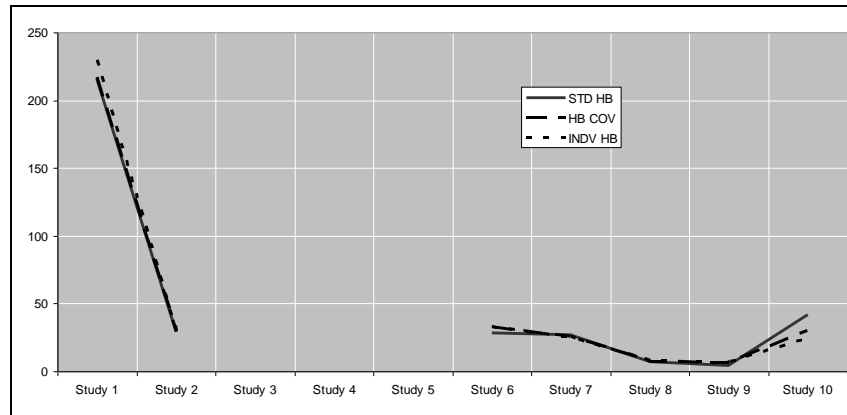
Graphic 3: MSE and Hit Rates of different simulation models

Especially the results of LC showed a controversial picture: In study 5 LC led to the worst MSE result while the Hit Rates (based on cluster members averages) scored best.

Overall there was no significant improvement in most of the studies through usage of covariates (either with natural segments or LC or Utility cluster based). Also, other estimation methods like LC or single HB segment estimation did not exceed the results achieved with standard HB in a significant way.

Holdout task results are mostly used for measurement of MSE and individual Hit Rates. As we had real market data for 7 of the 10 studies we used this calculate the MSE of the non-calibrated simulations (no correction for distribution and other external effects) as the ultimate proof of validity.

With exception of study one, which has a small and fragmented sample as well as a simplified attribute/level model, there were neither real differences between real market data and data from our studies nor significant differences between STD HB, HB COV and INDV HB as the graphic below shows:



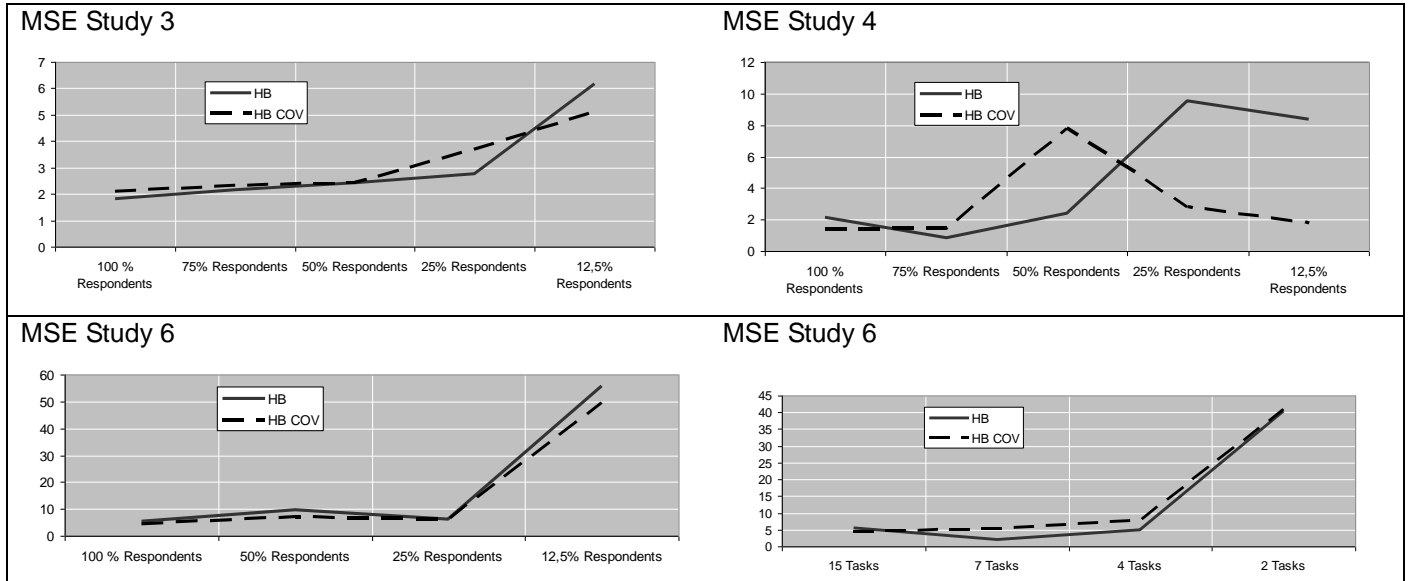
Graphic 4: MSE of not calibrated simulations against real Market Data

Our first conclusion is that studies which are set up correctly and have large enough sample size in all subgroups don't show better estimates when using covariates.

ESTIMATIONS WITH WEAKENED DATA

Based on our experience with the different estimation models and in order to simulate sparse data sets or poorly designed studies, we analyzed the effect of covariates on weakened data. In 3 of the ten studies we reduced the number of respondents stepwise randomly from 100% to 25%. The next try was to reduce number of tasks stepwise from 15 to 2 tasks by deleting the later tasks from the interview process.

To our surprise the reduction of "some amount" of information had no significant effect on the estimations. In two studies the MSE results remained on the same level even with only 50% of respondents, or as in study 6, with 25%.



Graphic 5: Simulation results with stepwise weakened data

Only in study 4 HB COV helped after reducing the sample by 50% or more. Otherwise there was no big difference between the HB with and without covariates. Furthermore once the information lack became too large (e.g 15% of sample) the error increased dramatically. However, again the Covariate was also not able to improve the simulation results.

The second conclusion is therefore that covariates are not a "first aid kit" for badly designed studies. Even though information could be reduced to some extent without damage, sample size must be retained to ensure representativeness.

As all of the previous analyses were based on the accuracy of prediction for the markets covered by the ten different studies, it was necessary to investigate how covariates could impact on the simulation of single segments of a study respectively of a market. For this purpose we selected two of the 10 studies and ran MSE and Hit Rate analysis within the natural market segments. As graphic 6 shows, there were no significant differences between the different estimation models. In study 4 there were three segments, and in each segment a different estimation model performed best. In the six segments of study 5 there was also no clear winner:




Graphic 6: Simulation results within market segments

As there was no significant effect in within-segment estimation between different variants of HB estimations to be observed, the third conclusion is that with disproportional sample structure (same/sufficient sample size for segment cells as there were in all ten studies) no improvement can be achieved by using Covariates.

ESTIMATIONS WITH PROPORTIONAL SEGMENTS

Study 6 had a quite large sample size. This was also caused by the disproportional sample structure: comparatively small segments were surveyed and analyzed with a sample size that was much higher than their representative market share, thus ensuring enough data for later choice estimation. Of course, these simulation results need weighting to the real segment share in a total market model. For further analysis we adapted the disproportional sample structure of study 6 (N=8,900 interviews) to the proportional market weights eliminating 3,300 interviews (N=5,600).

Introducing natural proportions into the 4 small segments, which then have insufficient sample size, led to generally worse estimates. In all 4 segments covariates were nevertheless able to improve the results. However, these were far away from the accuracy we observed with disproportional samples.



Graphic 7: Covariates in Proportional sample segments

The fourth conclusion is therefore that covariates do not allow for proportional sampling of small segment cells

TYPES OF COVARIATES AND IMPACT ON RESULTS

It is often stated that covariates work best when they add *new* information to the CBC data, and when the covariate information is strongly predictive of respondent preferences. Furthermore, it is stated that variables related to behavior and preferences will tend to be more valuable covariates than descriptive information such as demographics. Therefore we tested brand preference, past purchase, and available budget in some of our studies as potential candidates to be used as covariates. In general we could see that such added data did not result in real improvement of hit rates or MSE. If the additional information is chosen carefully it doesn't affect the HB estimation very much, but it also could degrade the results if the additional information is contrary to the original data.

Segmentation solutions based on cluster analysis of dozens of variables including preferences, attitudes, and psychographics could be valuable when introduced as categorical covariates. However, from a practitioner's point of view it is hard to know whether the estimation model will benefit. We saw in most of our studies that such complex covariates didn't improve the results compared to real market data. In nearly all of our cases, when there were changes in part-worths because of covariates, we were not able to explain the direction in which the covariates changed the results. Our attempts to use Latent Class segment membership as covariates led to small improvements of hit rates and MSE, but quite often also showed large changes in the resulting part worths. Covariates developed using only the choice data tend not to be helpful, and generally lead to over-fitting. One reason for this over fitting is that no new information from outside the CBC data is being used. The information that was already available within the CBC data was, in essence, being used twice.

In further analysis we tried to use combinations of several covariates in one model at a time. The results showed that it's generally not ideal to include several covariates without first confirming their potential usefulness by testing the distribution of the choice data compared to the additional variables. We learned that it is much better to focus on just a few covariates thereby adding relatively few columns to the covariate design matrix. One potential saving of parameters could be to treat a covariate as continuous rather than to categorize it as dummies. By using a continuous variable as a covariate one can save many parameters to be estimated without sacrificing much information. However, in many cases there is no such continuous information available. As with any multiple regression application one should carefully examine whether the covariates are influenced by multicollinearity when using more than one covariate in the HB model. But our observation was that in only one of ten studies were small improvements detected through adding combinations of covariates.

The more sparse a dataset is (e.g. relatively few choice tasks compared to the number of parameters to estimate or small samples), the more Bayesian shrinkage toward the pooled upper-level model can be expected. Therefore we examined whether covariates were most effective with sparse data sets. For those datasets with sufficient information at the individual level, the Bayesian shrinkage is already relatively small in standard HB, and covariates have a limited ability to improve the results of small sample cells. At the same time there is a risk that covariates have a negative impact on the results (i.e. MSE compared to real market data being larger). The ten studies we analyzed showed that the covariates could help to improve the results in some of our small (proportional) sample cells. However, the accuracy was still less than results from disproportional samples. Although we found that the HB shrinkage was reduced to some degree, the use of covariates did not provide the hit rates and preference share accuracy needed when communicating results to clients.

TO PUT IT IN A NUTSHELL

Our examination of covariates showed the following: The use of covariates is not really time saving. It neither results in shorter estimation time, nor is it fast and simple to use. The number of different estimation runs necessary to identify those covariates that improve or diminish the results required a lot of computational time and did not show any advantage against other techniques that could be used to reduce the Bayesian shrinkage.

The application of covariates reduced the precision of the estimates as often as it improved them. We found that the application of covariates is neither a fast nor an easy technique for everyday work. A lot of experience and analytical work on data structure is necessary in order to gain profound knowledge about the distribution of the data and to ensure that the applied covariates will really improve the results. The hypothesis that covariates increase the accuracy of estimations in regard to MSE, hit rates or real market data could not be confirmed in any of our ten studies analyzed.

The hypothesis that smaller samples or fewer tasks are needed when covariates are introduced could not be confirmed either. We tried systematically reducing sample sizes and numbers of tasks per respondent, in an attempt to see whether covariates helped when there was less information. But we saw that these manipulations had little effect on the quality of results, and the differences in quality occurred independently of the use of covariates. A rationale for this

phenomenon could be that in many cases we use too-long interviews and therefore get noise into our data by burdening the respondents (but this should be the topic of another paper).

Perhaps covariates could be used as a first aid kit if one observes that the data is sparse when estimating with standard HB or that there are too many parameters in the model to reach convergence in the estimates. Perhaps carefully chosen covariates adding additional exogenous information could help to improve the results in some cases. But one always should be aware that this improvement is only marginal compared to the results based on well-designed studies. Therefore we could conclude that covariates are not a “gold standard” for estimation. They could sometimes be helpful, but normally we would recommend the use of standard HB.

In most of our observations covariates in general did not improve results. In studies with large enough segment cells the covariate model converges towards same estimates as with standard HB (No influence of the covariate). Our “Gold Standard” from a practical point of view: Ensure sufficient sample size (disproportional segment cells), and use standard HB with enough iterations to assure convergence.

Covariates could improve results if we have already clearly defined clusters with groups of respondents with different multi-normal distributions on attributes in the data. Different densities in different regions of the data structure could also be a good indicator for use of covariates. We therefore suggest that one should first analyze the density structure of the common distribution of the choice data carefully and then decide whether or not to use covariates or other techniques to improve the results.

If there is a strong underlying cluster structure in the population, which was not taken into account in the sample planning and which can be identified and added as covariates to the HB estimation, this could help to improve the results. But other techniques like using proportional sampling or the weighting technique (Howell 2007) could solve the problem too. In our experience it is preferable not to add too much additional information at a time. This means trying each covariate within a single estimation before adding multiple covariates to your data.

Academics have shown that covariates can be superior in simulated environments. But in our re-analysis of ten studies, we show that using covariates can also be risky. In these studies we have had market data available to evaluate the effects of covariates, without which it would seem hard to evaluate the correct use of covariates. Covariates can in some cases improve estimates of parameters, but unfortunately not in the same amount than techniques like proportional sampling or alternative specific designs do.

REFERENCES

- Fuchs, S., Schwaiger M. “Disproportionate Samples in Hierarchical Bayes CBC Analysis” in: Decker, R., Lenz, H.-J. (Eds.) „Advances in Data Analysis“ Berlin, Heidelberg (Springer), pp 441-448.
- Howell, John (2007), “Respondent Weighting in HB,” Sawtooth Software Conference Proceedings, Sequim WA, pp 365-377.
- Lenk, P. J., DeSarbo, W. S., Green P. E. and Young, M. R. (1996) “Hierarchical Bayes Conjoint Analysis: Recovery of Partworth Heterogeneity from Reduced Experimental Designs,” *Marketing Science*, 15, 173-191.
- Orme, B.; Howell, J.(2009) “Application of Covariates within Sawtooth Software’s CBC/HB Program: Theory and Practical Example”, Sawtooth Software Technical Papers available at www.sawtoothsoftware.com
- Sawtooth Software, "More Evidence CBC Is Most Popular Conjoint Method," Sawtooth Solutions Newsletter, Summer 2003, available at www.sawtoothsoftware.com.
- Sawtooth Software (2005), “CBC/HB Technical Paper,” available at www.sawtoothsoftware.com.
- Sentis, Keith and Lihua Li, “One Size Fits All or Custom Tailored: Which HB Fits Better?” Sawtooth Software Conference Proceedings, Sequim, WA, pp 167-177.