Brand Imagery Measurement

Assessment of Current Practice and a New Approach

Brand Imagery research is an important and common component of market research programs. Traditional approaches, e.g., ratings scales, have serious limitations and may even sometimes be misleading.

Author: Paul Richard "Dick" McCullough

2013 Sawtooth Software Conference Proceedings,
October 2013, Dana Point, CA, USA

Brand Imagery Measurement

Assessment of Current Practice and a New Approach[1]

Executive Summary

Brand imagery research is an important and common component of market research programs. Traditional approaches, e.g., ratings scales, have serious limitations and may even sometimes be misleading.

MaxDiff scaling adequately addresses the major problems associated with traditional scaling methods but historically has had, within the context of brand imagery measurement, at least two serious limitations of its own. Until recently, MaxDiff scores were comparable only to items within the MaxDiff exercise. Traditional MaxDiff scores are relative, not absolute. Dual Response MaxDiff has substantially reduced this first problem but may have done so at the price of reintroducing scale usage bias. The second problem remains: MaxDiff exercises that span a reasonable number of brands and brand imagery statements often take too long to complete.

The purpose of this paper is to review the practice and limitations of traditional brand measurement techniques and to suggest a novel application of Dual Response MaxDiff that provides a superior brand imagery measurement methodology that increases inter-item discrimination and predictive validity and eliminates both brand halo and scale usage bias.

Introduction

Brand imagery research is an important and common component of most market research programs. Understanding the strengths and weaknesses of a brand, as well as its competitors, is fundamental to any marketing strategy. Ideally, any brand imagery analysis would not only include a brand profile, providing an accurate comparison across brands, attributes and respondents, but also an understanding of brand drivers or hot buttons.

Any brand imagery measurement methodology should, at a minimum, provide the following:

- Discrimination between attributes, for a given brand (inter-attribute comparisons)
- Discrimination between respondents or segments, for a given brand and attribute (inter-respondent comparisons)
- Good fitting choice or purchase interest model to identify brand drivers (predictive validity)

[1] The author wishes to thank Survey Sampling International for generously donating a portion of the sample used in this paper.

With traditional approaches to brand imagery measurement, there are typically three interdependent issues to address:

- Minimal variance across items, ie, flat responses
- Brand halo
- Scale usage bias

Resulting data are typically non-discriminating, highly correlated and potentially misleading. With high collinearity, regression coefficients may actually have reversed signs, leading to absurd conclusions, e.g., lower quality increases purchase interest.

While scale usage bias may theoretically be removed via modeling, there is reason to suspect any analytic attempt to remove brand halo since brand halo and real brand perceptions are typically confounded. That is, it is difficult to know whether a respondent's high rating of Brand A on perceived quality, for example, is due to brand halo, scale usage bias or actual perception.

Thus, the ideal brand imagery measurement technique will exclude brand halo at the data collection stage rather than attempt to correct for it at the analytic stage. Similarly, the ideal brand imagery measurement technique will eliminate scale usage bias at the data collection stage as well.

While the problems with traditional measurement techniques are well known, they continue to be widely used in practice. Familiarity and simplicity are, no doubt, appealing benefits of these techniques. Among the various methods used historically, the literature suggests that comparative scales may be slightly superior. An example of a comparative scale is below:

	Much more than other brands	Somewhat more than other brands	Neither more nor less than other brands	Somewhat less than other brands	Much less than other brands
Brand A makes high quality products	☐	☐	☐	☐	☐

Some alternative techniques have also garnered attention: MaxDiff scaling, method of paired comparisons and q-sort. With the exception of Dual Response MaxDiff (DR MD), these techniques all involve relative measures rather than absolute.

MaxDiff scaling, MPC and Q-sort all are scale-free (no scale usage bias), potentially have no brand halo[2] and demonstrate more discriminating power than more traditional measuring techniques.

MPC is a special case of MaxDiff; as it has been shown to be slightly less effective it will not be further discussed separately.

[2] These techniques do not contain brand halo effects if and only if the brand imagery measures are collected for each brand separately rather than pooled.

With MaxDiff scaling, the respondent is shown a random subset of items and asked to pick which he/she most agrees with and which he/she least agrees with. The respondent is then shown several more subsets of items. A typical MaxDiff question is shown below:

Traditional MaxDiff [3]

Thinking about the Sony brand, which statement do you Most associate with that brand, and which statement do you Least associate with that brand.

Most	SONY	Least
○	Available in stores near me	○
○	Makes products geared towards women	○
○	Makes products that children can use	○
○	Is trust-worthy	○
○	Expensive	○

With Q-sorting, the respondent is asked to place into a series of "buckets" a set of items, or brand image attributes, from best describes the brand to least describes the brand. The number of items in each bucket roughly approximates a normal distribution. Thus, for 25 items, the number of items per bucket might be:

First bucket	1 item
Second bucket	2 items
Third bucket	5 items
Fourth bucket	9 items
Fifth bucket	5 items
Sixth bucket	2 items
Seventh bucket	1 item

MaxDiff and q-sorting adequately address two of the major issues surrounding monadic scales, inter-attribute comparisons and predictive validity, but due to their relative structure do not allow inter-brand comparisons. That is, MaxDiff and q-sorting will determine which brand imagery statements have higher or lower scores than other brand imagery statements for a given brand but can't determine which brand has a higher score than any other brand on any given statement. Some would argue that MaxDiff scaling also does not allow inter-respondent comparisons due to the scale factor. Additionally, as a practical matter, both techniques currently accommodate fewer brands and/or attributes than traditional techniques.

Both MaxDiff scaling and Q-sorting take much longer to field than other data collection techniques and are not comparable across studies with different brand and/or attribute sets. Q-sorting takes less time to complete than MaxDiff and is somewhat less discriminating.

As mentioned earlier, MaxDiff can be made comparable across studies by incorporating the Dual Response version of MaxDiff, which allows the estimation of an absolute reference point. This reference

[3] The form of Max/Diff scaling used in brand imagery measurement is referred to as Brand-Anchored Max/Diff (BA MD)

point may come at a price. The inclusion of an anchor point in MaxDiff exercises may reintroduce scale usage bias into the data set.

However, for q-sorting, there is currently no known approach to establish an absolute reference point. For that reason, q-sorting, for the purposes of this paper, is eliminated as a potential solution to the brand measurement problem.

Also, for both MaxDiff and q-sorting the issue of data collection would need to be addressed. As noted earlier, to remove brand halo from either a MaxDiff-based or q-sort-based brand measurement exercise, it will be necessary to collect brand imagery data on each brand separately, referred to here as brand-anchored MaxDiff. If the brands are pooled in the exercise, brand halo would remain. Thus, there is the very real challenge of designing the survey in such a way as to collect an adequate amount of information to accurately assess brand imagery at the disaggregate level without overburdening the respondent.

Although one could estimate an aggregate level choice model to estimate brand ratings, that approach is not considered viable here because disaggregate brand ratings data are the current standard. Aggregate estimates would yield neither familiar nor practical data. Specifically, without disaggregate data, common cross tabs of brand ratings would be impossible as would the more advanced predictive model-based analyses.

A New Approach

Brand-anchored MaxDiff, with the exception of being too lengthy to be practical, appears to solve, or at least substantially mitigate, most of the major issues with traditional methods of brand imagery measurement. The approach outlined below attempts to minimize the survey length of brand-modified MaxDiff by increasing the efficiency of two separate components of the research process:

- Survey instrument design
- Utility estimation

Survey Instrument

A new MaxDiff question format, referred to here as modified Brand-anchored MaxDiff, accommodates more brands and attributes than the standard design. The format of the modified Brand-anchored MaxDiff used in Image MD© is illustrated below:

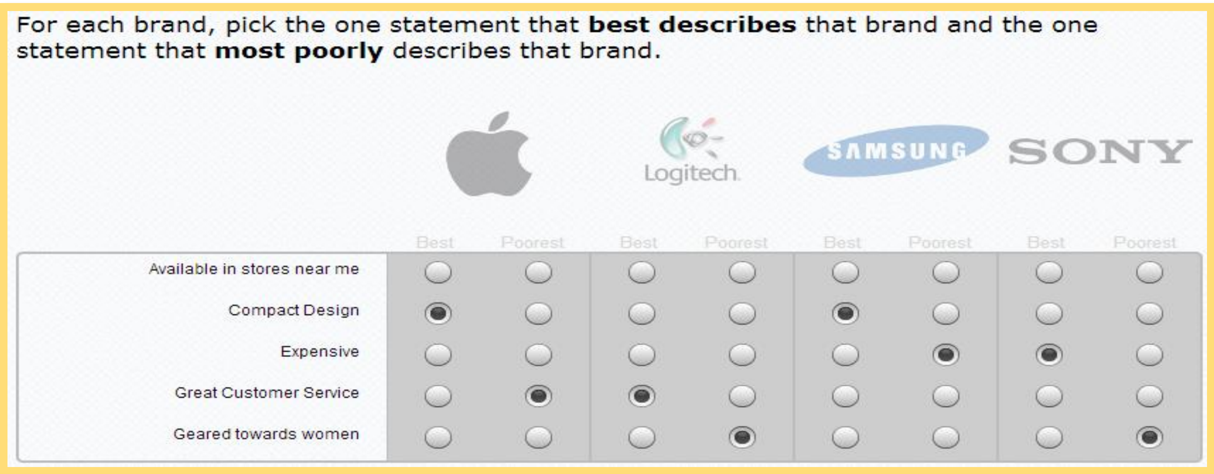

To accommodate the Dual Response form of MaxDiff, a Direct Binary Response question is asked prior to the MBA MD task set[4]:

To address the potential scale usage bias of MaxDiff exercises with Direct Binary Response, a negative Direct Binary Response question, eg, *For each brand listed below, please check all the attributes that you feel strongly do not describe the brand, is also included.*[5] As an additional attempt to mitigate scale usage bias, the negative Direct Binary Response was asked in a slightly different way for half the sample. Half the sample were asked the negative Direct Binary Response question as above. The other half were asked a similar question except that respondents were required to check as many negative items as they had check positive. The first approach is referred to here as unconstrained negative Direct Binary Response and the second is referred to as constrained negative Direct Binary Response.

In summary, Image MD© consists of an innovative MaxDiff exercise and two direct binary response questions, as shown below:

[4] This approach to Anchored Max/Diff was demonstrated to be faster to execute than the traditional Dual Response format (Lattery 2010).

[5] Johnson and Fuller (2012) note that Direct Binary Response yields a different threshold than traditional Dual Response. By collecting both positive and negative Direct Binary Response data, we will explore ways to mitigate this effect.

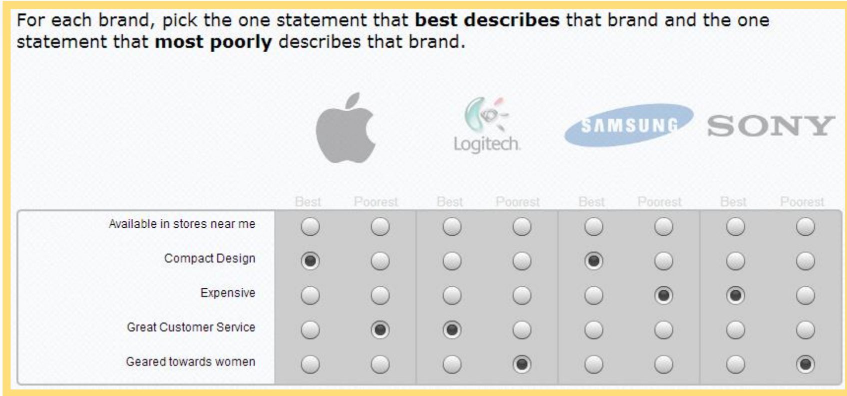

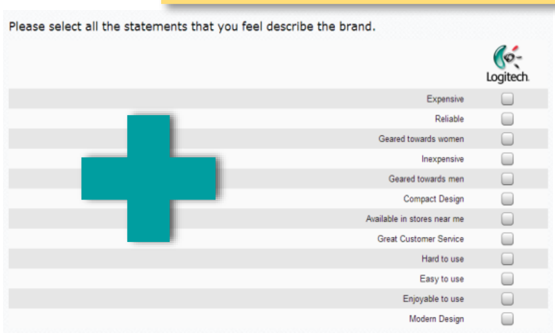

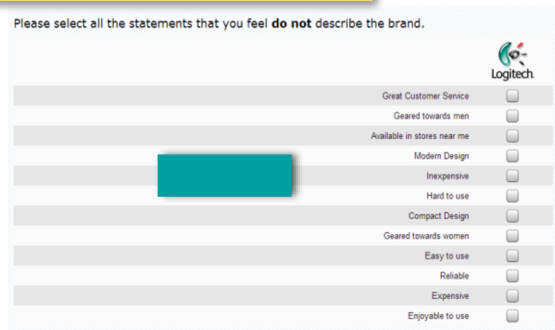

It is possible, in an online survey, to further increase data collection efficiency with the use of some imaginative programming. We have developed an animated way to display Image MD© tasks which can be viewed at www.macroinc.com (Research Techniques tab, MaxDiff Item Scaling).

Thus, the final form of the Image MD© brand measurement technique can be described as Animated Modified Brand-Anchored MaxDiff Scaling with both Positive and Negative Direct Binary Response.

Utility Estimation

Further, an exploration was conducted to reduce the number of tasks seen by any one respondent and still retain sufficiently accurate disaggregate brand measurement data. MaxDiff utilities were estimated using a Latent Class Choice Model (LCCM) and using a hierarchical Bayes model (HB). By pooling data across similarly behaving respondents (in the LCCM), we hoped to substantially reduce the number of MaxDiff tasks per respondent. This approach may be further enhanced by the careful use of covariates. Another approach that may require fewer MaxDiff tasks per person is to incorporate covariates in the upper model of an HB model or running separate HB models for segments defined by some covariate.

To summarize, the proposed approach consists of:

- Animated Modified Brand-Anchored MaxDiff Exercise
- With Direct Binary Responses (both positive and negative)
- Analytic-derived parsimony:
  - Latent Class Choice Model:
    - Estimate disaggregate MaxDiff utilities
    - Use of covariates to enhance LCCM accuracy
  - Hierarchical Bayes:
    - HB with covariates in upper model

- Separate HB runs for covariate-defined segments
- Adjusted priors[6]

Research Objectives

The objectives, then, of this paper are:

- To compare this new data collection approach, Animated Modified Brand-Anchored MaxDiff with Direct Binary Response, to a traditional approach using monadic rating scales
- To compare the positive Direct Binary Response and the combined positive and negative Direct Binary Response
- To confirm that Animated Modified Brand-Anchored MaxDiff with Direct Binary Response eliminates brand halo
- To explore ways to include an anchor point without reintroducing scale usage bias
- To explore utility estimation accuracy of LCCM and HB using a reduced set of MaxDiff tasks
- To explore the efficacy of various potential covariates in LCCM and HB

Study Design

A two cell design was employed: Traditional brand ratings scales in one cell and the new MaxDiff approach in the other. Both cells were identical except in the method that brand imagery data were collected:

- Traditional brand ratings scales
  - Three brands, each respondent seeing all three brands
  - 12 brand imagery statements
- Animated Modified Brand-Anchored MaxDiff with Direct Binary Response
  - Three brands, each respondent seeing all three brands
  - 12 brand imagery statements
  - Positive and negative Direct Binary Response questions

Cells sizes were:

- Monadic ratings cell – n = 436
- Modified MaxDiff – n = 2,605
  - Unconstrained negative DBR – n = 1,324
  - Constrained negative DBR – n = 1,281

The larger sample size for the second cell was intended so that attempts to reduce the minimum number of choice tasks via LCCM and/or HB could be fully explored.

Both cells contained:

- Brand imagery measurement (ratings or MaxDiff)
- Brand affinity measures

[6] McCullough (2009) demonstrates that tuning HB model priors can improve hit rates in sparse data sets.

- Demographics
- Holdout attribute rankings data

Results

Brand Halo

We check for brand halo using confirmatory factor analysis, building a latent factor to capture any brand halo effect. If the brand halo exists, the brand halo latent factor will positively influence scores on all items. We observed a clear brand halo effect among the ratings scale data, as expected. The unanchored MaxDiff data showed no evidence of the effect, also as expected. The positive direct binary response reintroduced the brand halo effect to the MaxDiff ratings at least as strong as the ratings scale data. This was not expected. However, the effect seems to be totally eliminated with the inclusion of either the constrained or unconstrained negative direct binary question.

Brand Halo Confirmatory Factor Analytic Structure

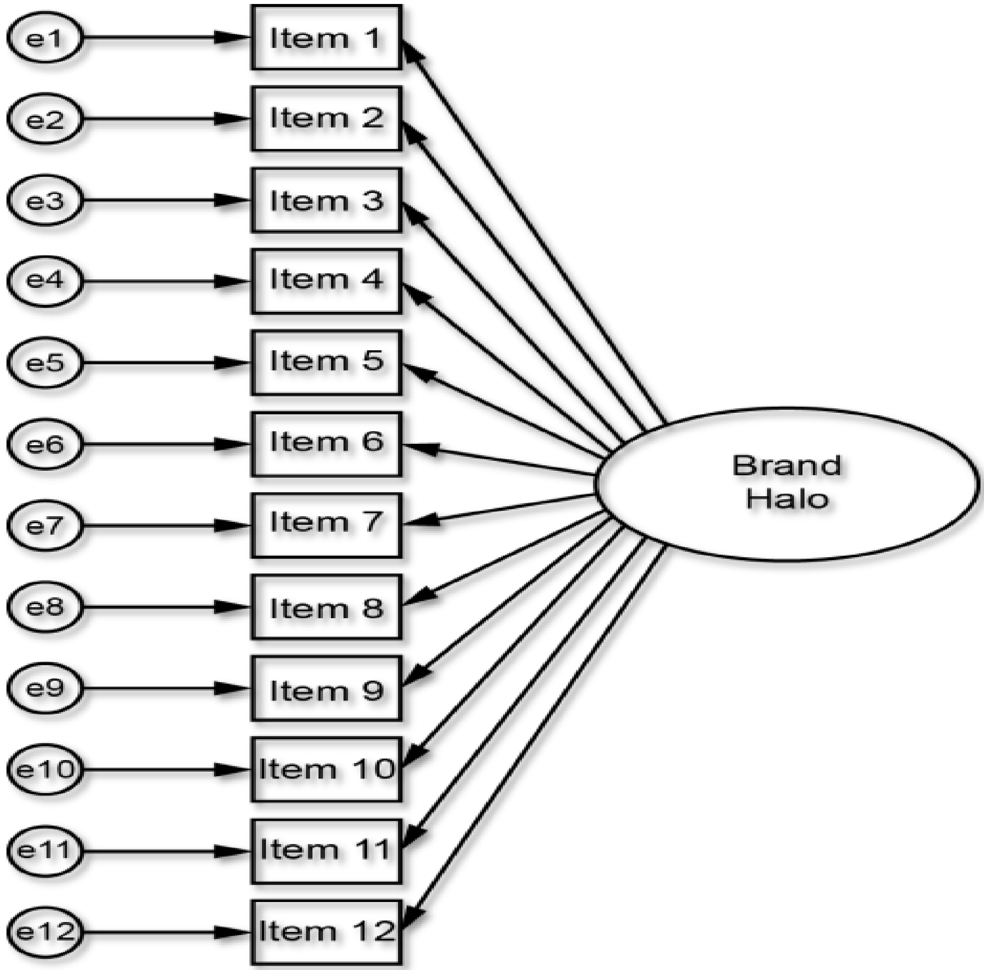

Brand Halo Latent	Ratings		No DBR		Positive DBR		Unconstrained Negative DBR		Constrained Negative DBR	
	Std Beta	Prob	Std Beta	Prob	Std Beta	Prob	Std Beta	Prob	Std Beta	Prob
Item 1	0.85	***	-0.14	***	0.90	***	0.44	***	0.27	***
Item 2	0.84	***	-0.38	***	0.78	***	-0.56	***	-0.72	***
Item 3	0.90	***	-0.20	***	0.95	***	0.42	***	0.32	***
Item 4	0.86	***	0.10	***	0.90	***	0.30	***	0.16	***
Item 5	0.77	***	-0.68	***	0.88	***	0.03	0.25	0.01	0.78
Item 6	0.85	***	-0.82	***	0.87	***	-0.21	***	-0.24	***
Item 7	0.83	***	0.69	***	0.83	***	0.42	***	0.20	***
Item 8	0.82	***	0.24	***	0.75	***	0.01	0.87	-0.23	***
Item 9	0.88	***	0.58	***	0.90	***	0.77	***	0.62	***
Item 10	0.87	***	0.42	***	0.94	***	0.86	***	0.90	***
Item 11	0.77	***	-0.05	0.02	0.85	***	0.07	0.02	-0.12	***
Item 12	0.88	na	0.26	na	0.91	na	0.69	na	0.53	na

Scale Usage

As with our examination of brand halo, we use confirmatory factor analysis to check for the presence of a scale usage factor. We build in latent factors to capture brand halo per brand, and build another latent factor to capture a scale usage bias independent of brand. If a scale usage bias exists, the scale latent factor should load positively on all items for all brands.

Scale Usage Bias and Brand Halo Confirmatory Factor Analytic Structure

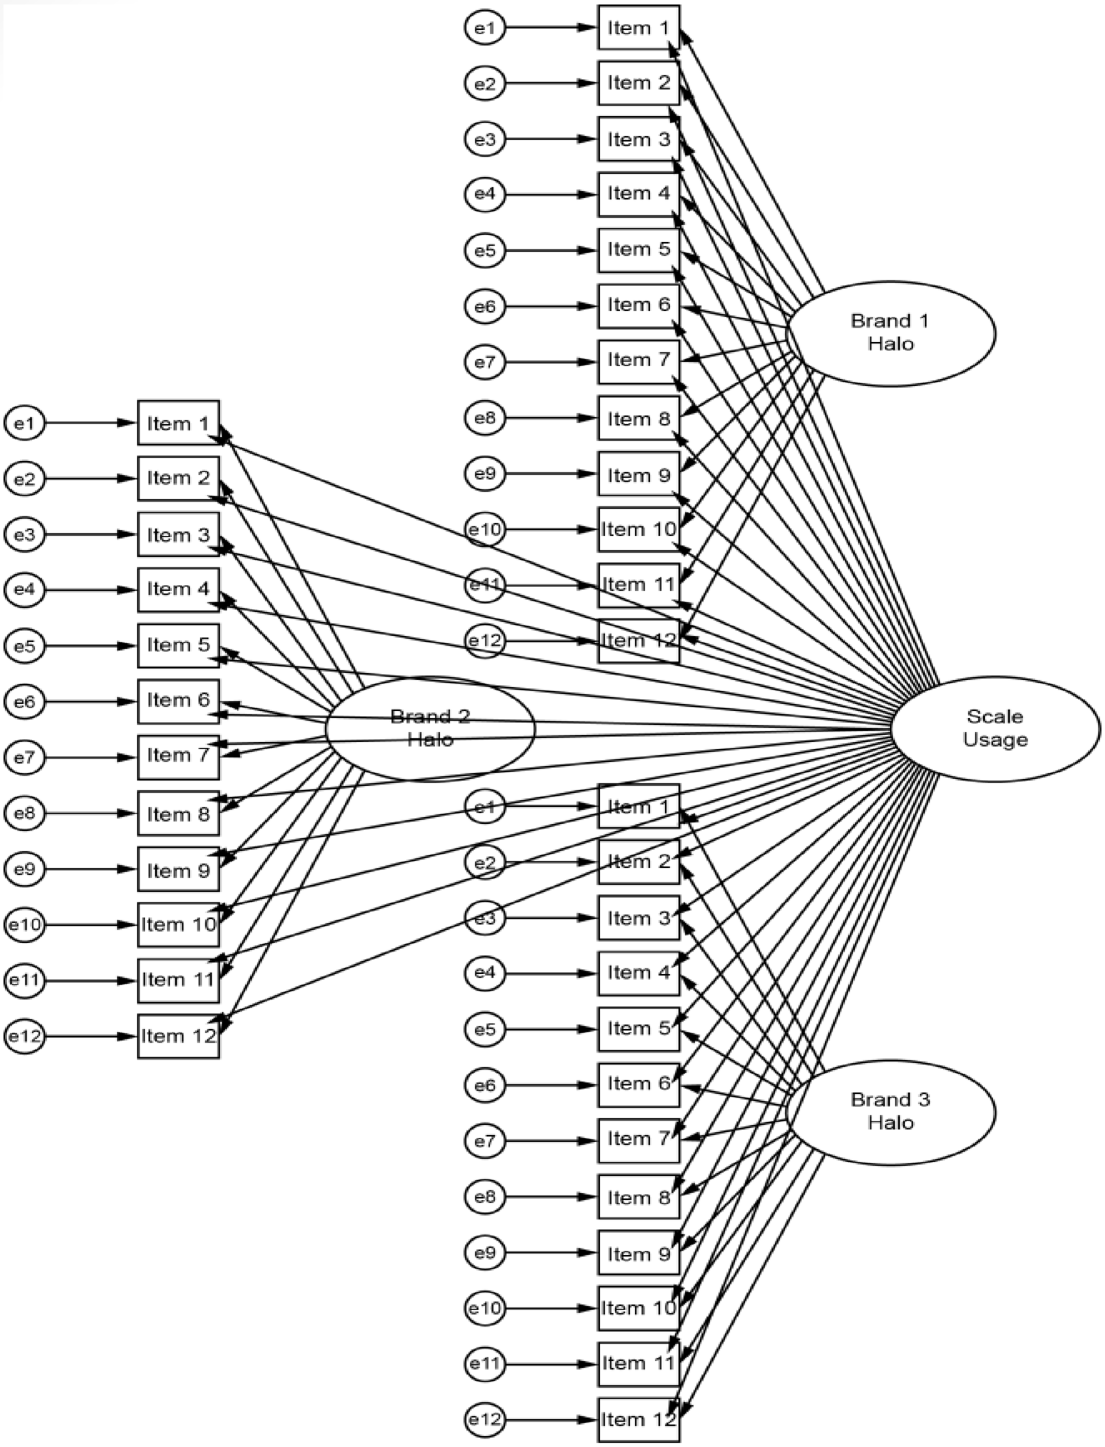

We observe an obvious scale usage effect with the ratings data, where the scale usage latent loads positively on all 36 items. Again, the MaxDiff with only positive direct binary response showing some indication of scale usage bias, even with all three brand halo latents simultaneously accounting for a great deal of collinearity. Traditional MaxDiff, and the two versions including positive and negative direct binary responses all show no evidence of a scale usage effect.

Scale Usage Latent	Ratings	No DBR	Positive DBR	Unconstrained Negative DBR	Constrained Negative DBR
Number of Negative Loadings	0	14	5	10	15
Number of Statistically Significant loadings	36	31	29	33	30

Predictive Validity

In the study design we included a holdout task which asked respondents to rank their top three item choices per brand, giving us a way to test the accuracy of the various ratings/utilities we collected. In the case of all MaxDiff data we compared the top three scoring items to the top three ranked holdout items per person, and computed the hit rate. This approach could not be directly applied to scale ratings data due to the frequency of flat responses (e.g. it is impossible to identify top three if all items were rated the same). For the ratings data we estimated hit rate using this approach: if the highest rated item from the holdout received the highest ratings score which was shared by n items, we added 1/n to the hit rate. Similarly, the second and third highest ranked holdout items received an adjusted hit point if those items were among the top 3 rated items.

We observe that each of the MaxDiff data sets vastly outperformed ratings scale data, which performed roughly the same as randomly guessing the top three ranked items.

Hit Rates	Random Numbers	Ratings	No DBR	Positive DBR	Unconstrained Negative DBR	Constrained Negative DBR
1 of 1	8%	14%	27%	28%	27%	26%
(1 or 2) of 2	32%	30%	62%	64%	62%	65%
(1, 2 or 3) of 3	61%	51%	86%	87%	86%	88%

Inter-item discrimination

Glancing visually at the resulting item scores, we can see that each of the MaxDiff versions show greater inter-item discrimination, and among those, both negative direct binary versions bring the lower performing brand closer to the other two brands.

Ratings Scales

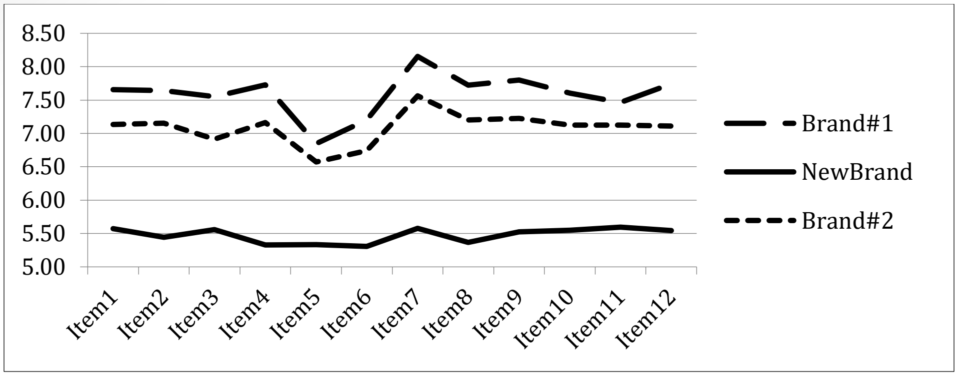

MaxDiff with Positive DBR

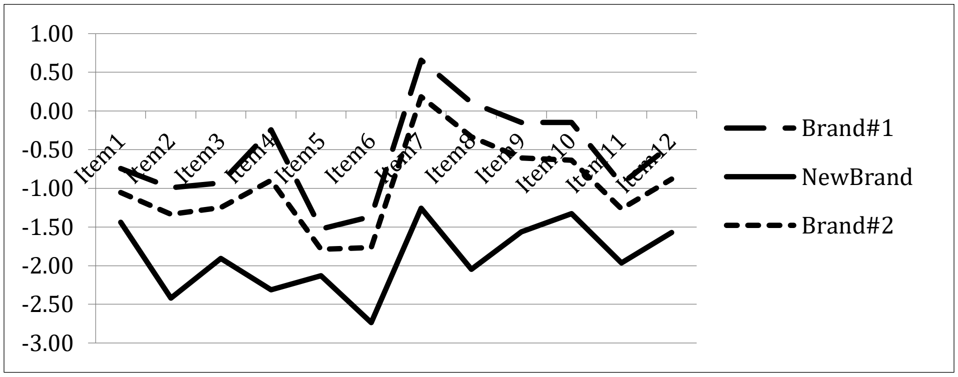

MaxDiff with Positive DBR & Constrained Negative DBR

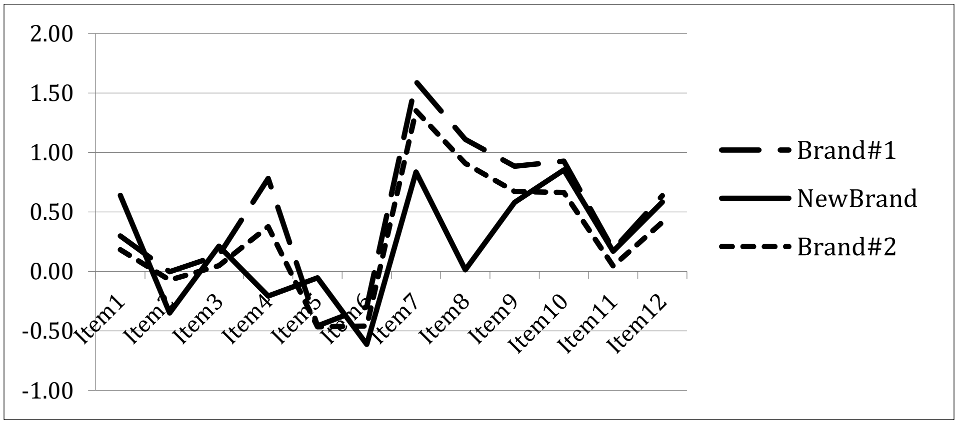

MaxDiff with Positive DBR & Unconstrained Negative DBR

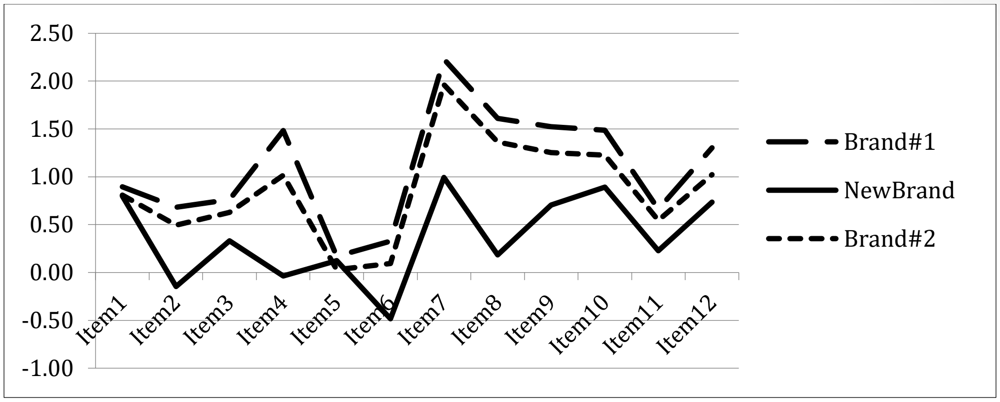

To confirm, we considered how many statistically significant differences between statements could be observed within each brand per data collection method. The ratings scale data yielded fewest statistically significant differences across items, while the MaxDiff with positive and unconstrained negative direct binary responses yielded the most. Traditional MaxDiff and MaxDiff with positive and constrained negative direct binary responses also performed very well, while the MaxDiff with only positive direct binary performed much better than ratings scale data, but clearly not as well as the remaining three MaxDiff methods.

Average number of statistically significant differences across 12 items	Ratings	No DBR	Positive DBR	Unconstrained Negative DBR	Constrained Negative DBR
Brand#1	1.75	4.46	3.9	4.3	4.68
New Brand	0	4.28	3.16	4.25	4.5
Brand#2	1	4.69	3.78	4.48	4.7

Completion Metrics

With using a more sophisticated data collection method come a few costs in respondent burden. It took respondents much longer to complete any of the MaxDiff exercises than it took them to complete the simple ratings scales. The dropout rate during the brand imagery section of the survey (measured as percentage of respondents who began that section but failed to finish it) was also much higher among the MaxDiff versions. Though on the plus side for the MaxDiff versions, when preparing the data for analysis we were forced to drop far fewer respondents due to flat-lining.

	Ratings	All MaxDiff Versions
Brand Image Measurement Time (Minutes)	1.7	6
Incompletion Rate	9%	31%
Post-field drop rate	32%	4%

Exploration to Reduce Number of Tasks Necessary

We find these results to be generally encouraging, but would like to explore if anything can be done to reduce the increased respondent burden and dropout rates. Can we reduce the number of tasks each respondent is shown, without compromising the predictive validity of the estimated utilities? To find out, we estimated disaggregate utilities using two different estimation methods (Latent Class and Hierarchical Bayes), varying the numbers of tasks, and using certain additional tools to bolster the quality of the data (using covariates, or adjusting priors, etc.).

We continued only with the two MaxDiff methods with both positive and negative direct binary responses, as those two methods proved best in our analysis. All estimation routines were run for both the unconstrained and constrained versions, allowing us to further compare these two methods.

Our chosen covariates included home ownership (rent vs. own), gender, purchase likelihood for the brand we were researching, and a few others. Including these covariates when estimating utilities in HB should yield better individual results by allowing the software to make more educated estimates based on respondents' like peers.

With covariates in place, utilities were estimated using data from 8 (full sample), 4, and 2 MaxDiff tasks, and hit rates were computed for each run. We were surprised to discover that using only 2 tasks yielded only slightly less accuracy than using all 8 tasks. And in all cases, hit rates seem to be mostly maintained despite decreased data.

Using Latent Class the utilities were estimated again using these same 6 data sub-sets. As with HB, reducing the number of tasks used to estimate the utilities had minimal effect on the hit rates. It is worth noting here, that when using all 8 MaxDiff tasks latent class noticeably underperforms hierarchical bayes, but this disparity decreases as tasks are dropped.

Various Task Hit Rates		Unconstrained Negative DBR			Constrained Negative DBR		
		8 Tasks	4 Tasks	2 Tasks	8 Tasks	4 Tasks	2 Tasks
HB	1 of 1	27%	21%	20%	26%	24%	22%
	(1 or 2) of 2	62%	59%	58%	65%	61%	59%
	(1, 2 or 3) of 3	86%	85%	82%	88%	86%	85%
LC	1 of 1	19%	20%	19%	21%	21%	22%
	(1 or 2) of 2	54%	57%	56%	61%	59%	56%
	(1, 2 or 3) of 3	81%	82%	83%	84%	84%	82%

In estimating utilities in hierarchical Bayes, it is possible to adjust the Prior degrees of freedom and the Prior variance. Generally speaking, adjusting these values allows the researcher to change the emphasis placed on the upper level model. In dealing with sparse data sets, adjusting these values may lead to more robust individual utility estimates.

Utilities were estimated with data from 4 tasks, and with Prior degrees of freedom from 2 to 1000 (default is 5), and Prior variance from 0.5 to 10 (default is 2). Hit rates were examined at various points on these ranges, and compared to the default settings. After considering dozens of non-default configurations we observed essentially zero change in hit rates.

At this point it seemed that there was nothing that could diminish the quality of these utilities, which was a suspicious finding. In searching for a possible explanation, we hypothesized that these data simply have very little heterogeneity. The category of product being researched is not emotionally engaging (light bulbs), and the brands being studied are not very differentiated. To test this hypothesis, an additional utility estimation was performed, using only data from 2 tasks, and with a drastically reduced sample size of 105. Hit rates were computed for the low sample run both at the disaggregate level, that is using unique individual utilities, and then again with each respondents utilities equal to the average of the sample (constant utilities).

Unconstrained Negative DBR

	Random Choices	HB 8 Tasks N=1,324	HB 2 Tasks N=105	HB 2 Tasks N=105 Constant Utils
1 of 1	8%	27%	22%	25%
(1 or 2) of 2	32%	62%	59%	61%
(1, 2 or 3) of 3	61%	86%	82%	82%

These results seem to suggest that there is very little heterogeneity for our models to capture in this particular data; explaining why even low task utility estimates yield fairly high hit rates. Unfortunately, this means what we cannot say whether we can reduce survey length of this new approach by reducing the number of tasks needed for estimation.

Summary of Results

	Ratings	No DBR	Positive DBR	Unconstrained Negative DBR	Constrained Negative DBR
Provides Absolute Reference Point	No	No	Yes	Yes	Yes
Brand Halo	Yes	No	Yes	No	No
Scale Usage Bias	Yes	No	Yes	No	No
Inter-Item Discrimination	Very Low	High	Fairly High	High	High
Predictive Validity	Very Low	High	High	High	High
Complete Time	Fast	Slow	Slow	Slow	Slow
Dropout Rate	Low	High	High	High	High
Post-Field Drop Rate	High	Low	Low	Low	Low

Conclusions

The form of MaxDiff referred to here as Animated Modified Brand-Anchored MaxDiff Scaling with both Positive and Negative Direct Binary Response is superior to rating scales for measuring brand imagery:

- Better inter-item discrimination
- Better predictive validity
- Elimination of brand halo
- Elimination of scale usage bias
- Fewer invalid completes

Using positive DBR alone to estimate MaxDiff utilities reintroduces brand halo and possibly scale usage bias. Positive DBR combined with some form of negative DBR to estimate MaxDiff utilities eliminates both brand halo and scale usage bias. Utilities estimated with Positive DBR have slightly weaker inter-item discrimination than utilities estimated with Negative DBR.

The implication to these findings regarding DBR is that perhaps MaxDiff, if anchored, should always incorporate both positive and negative DBR since positive DBR alone produces highly correlated MaxDiff utilities with less inter-item discrimination.

Another, more direct implication, is that Brand-Anchored MaxDiff with both positive and negative DBR is superior to Brand-Anchored MaxDiff with only positive DBR for measuring brand imagery.

Animated Modified Brand-Anchored MaxDiff Scaling with both Positive and Negative Direct Binary Response takes longer to administer and has higher incompletion rates, however, and further work needs to be done to make the data collection and utility estimation procedures more efficient.

References

Bacon, L., Lenk, P., Seryakova, K., and Veccia, E. (2007), "Making MaxDiff more informative: statistical data fusion by way of latent variable modeling," 2007 Sawtooth Software Conference Proceedings, Santa Rosa, CA

Bacon, L., Lenk, P., Seryakova, K., and Veccia, E. (2008), "Comparing Apples to Oranges," Marketing Research Magazine, Spring 2008

Bochenholt, U. (2004), "Comparative judgements as an alternative to ratings: Identifying the scale origin," Psychological Methods

Chrzan, Keith and Natalia Golovashkina (2006), "An Empirical Test of Six Stated Importance Measures," International Journal of Marketing Research

Chrzan, Keith and Doug Malcom (2007), "An Empirical Test of Alternative Brand Systems," 2007 Sawtooth Software Conference Proceedings, Santa Rosa, CA

Chrzan, Keith and Jeremy Griffiths (2005), "An Empirical Test of Brand-Anchored Maximum Difference Scaling," 2005 Design and Innovations Conference, Berlin

Cohen, Steven H. (2003), "Maximum Difference Scaling: Improved Measures of Importance and Preference for Segmentation," Sawtooth Software Research Paper Series

Dillon, William R., Thomas J. Madden, Amna Kirmani and Soumen Mukherjee (2001), "Understanding What's in a Brand Rating: A Model for Assessing Brand and Attribute Effects and Their Relationship to Brand Equity," JMR© 2000 / MACRO Consulting, Inc.

Hendrix, Phil, and Drucker, Stuart (2007), "Alternative Approaches to Maxdiff With Large Sets of Disparate Items-Augmented and Tailored Maxdiff, 2007 Sawtooth Software Conference Proceedings, Santa Rosa, CA

Horne, Jack, Bob Rayner, Reg Baker and Silvo Lenart (2012), "Continued Investigation Into the Role of the 'Anchor' in MaxDiff and Related Tradeoff Exercises", 2012 Sawtooth Software Conference Proceedings, Orlando, FL

Johnson, Paul, and Brent Fuller (2012), "Optimizing Pricing of Mobile Apps with Multiple Thresholds in Anchored MaxDiff", 2012 Sawtooth Software Conference Proceedings, Orlando, FL

Lattery, Kevin (2010), "Anchoring Maximum Difference Scaling Against a Threshold-Dual Response and Direct Binary Responses", 2010 Sawtooth Software Conference Proceedings, Newport Beach, CA

Louviere, J.J., Marley, A.A.J., Flynn, T., Pihlens, D. (2009), "Best-Worst Scaling: Theory, Methods and Applications," CenSoc: forthcoming.

Magidson, J., and Vermunt, J.K. (2007a), "Use of a random intercept in latent class regression models to remove response level effects in ratings data. Bulletin of the International Statistical Institute, 56th Session, paper #1604, 1-4. ISI 2007: Lisboa, Portugal

Magidson, J., and Vermunt, J.K. (2007b), "Removing the scale factor confound in multinomial logit choice models to obtain better estimates of preference," 2007 Sawtooth Software Conference Proceedings, Santa Rosa, CA

McCullough, Paul Richard (2009), "Comparing Hierarchical Bayes and Latent Class Choice: Practical Issues for Sparse Data Sets", 2009 Sawtooth Software Conference Proceedings, Delray Beach, FL

Vermunt and Magidson (2008), LG Syntax User's Guide: Manual for Latent GOLD Choice 4.5 Syntax Module

Wirth, Ralph, and Wolfrath, Annette (2012), "Using MaxDiff for Evaluating Very Large Sets of Items", 2012 Sawtooth Software Conference Proceedings, Orlando, FL

Published in the 2013 SAWTOOTH SOFTWARE CONFERENCE PROCEEDINGS, October 2013,
Dana Point, CA, USA

We are an independent ***marketing research consulting firm*** *dedicated to helping you make the most informed, insightful marketing decisions possible. We specialize in technology, consumer, and new product research, and are well recognized for our* ***State-of-the-Art Research*** *techniques.*

Ultimately, we provide more than just technical expertise. We focus on developing ***pragmatic solutions*** *that will have a positive impact on the profitability of our clients.*

CONTACT US:

Telephone: 650-823-3042

General Inquiries:
info@macroinc.com

Advanced Analysis Inquiries:
analysis@macroinc.com

richard@macroinc.com

www.macroinc.com