



# Sawtooth Software

## *RESEARCH PAPER SERIES*

### An Empirical Comparison of ACA and Full Profile Judgments

Joel C. Huber, Duke University

Dick R. Wittink, Cornell University

John A. Fiedler, Populus, Inc.

and

Richard L. Miller, Consumer Pulse, Inc.

1991

# AN EMPIRICAL COMPARISON OF ACA AND FULL PROFILE JUDGMENTS

*Joel C. Huber*  
Duke University  
*Dick R. Wittink*  
Cornell University  
*John A. Fiedler*  
POPULUS, Inc.  
*Richard L. Miller*  
Consumer Pulse, Inc.

Published in Sawtooth Software Conference Proceedings, 1991

## INTRODUCTION

Only a few years ago, Johnson (1987) introduced Adaptive Conjoint Analysis (ACA) to the market research community. This innovative approach to the quantification of consumers' preference structures has become very popular (Green et al. 1991). One advantage of ACA, relative to other approaches, is that it combines the design of the conjoint tasks, data collection, data analysis and market simulation in one piece of software. Perhaps the most interesting aspect of ACA is that the use of the personal computer allows ACA to adapt the conjoint task to the individual respondent. The procedure can be considered to consist of two stages. In the first stage, the respondent indicates, among other things, the (relative) importance of each of the attributes defined for the study. In the second stage, answers to conjoint questions are elicited. These questions are defined based on information gathered in the first stage (and based on answers to previous conjoint questions). In this manner, the conjoint tasks can elicit responses to questions designed to reduce the uncertainty in the respondent's estimated parameters. Although ACA can be used for any problem that requires a quantification of preference structures, it should become increasingly attractive (relative to alternative methods) as 1) the number of attributes to be included in the study increases, and 2) respondent heterogeneity in the attributes' relevance increases.

Given the availability of several alternative procedures for preference (and choice) measurement, it is of interest to evaluate and compare their performances. Green and Srinivasan (1990) recommend full profile conjoint analysis if no more than approximately six attributes are involved in a study. For a somewhat larger number of attributes, they suggest that tradeoff matrices could be used (although they also favor bridging designs with full profiles). Only when the number of attributes reaches ten or more do they recommend the use of self-explicated data and methods involving a combination of self-explicated and conjoint data (such as ACA).

Currently there is very little information available on the (empirical) performance of alternative methods. MacBride and Johnson (1979) obtained higher first choice predictive validity for a forerunner of ACA relative to the tradeoff matrix approach, although the difference was not statistically significant. Finkbeiner and Platz (1986) compared ACA with the full profile method in a study involving six monotone attributes. They obtained roughly comparable predictive

validities.

Agarwal and Green (1989) had respondents go through ACA, a separate self-explicated task and two full profile rating tasks. They also used six monotone attributes. Although Agarwal and Green obtained better predictions from the self-explicated data (than ACA) on the full profile holdout data, Johnson (1991) expresses substantial concerns about confounds.

In this paper, we focus on ACA and full profile as two alternative methods for quantifying preference structures. We mention unique characteristics of each method that can be used to suggest application contexts for which a given method is especially suitable. This is followed by a description of our study to obtain empirical results on the relative performance of the two methods. We show predictive validity results and end with conclusions and suggestions for further research.

### **ACA Versus Full Profile**

Given that ACA and full profile appear to be the most popular approaches for quantifying consumers' preference structures, it is of interest to ask more generally how these two methods compare. One may argue that the two approaches are so different that each has its own application areas. On the other hand, there are many cases where either approach can be used. It is important then to provide both conceptual (theoretical) and empirical comparisons. For a detailed description of ACA, see Johnson (1987).

We provide in Table 1 a comparison of some distinct and relevant characteristics of the two approaches. The descriptions in this table are a convenient summary of the essential differences. We believe it is useful to mention that there is potential within ACA to impose constraints on the estimated partworths. For full profiles, Srinivasan *et al.* (1983) show that the predictive validity can be improved if homogeneous *a priori* constraints on the order of individual attributes' partworths are imposed for all respondents. In ACA, such constraints can easily be introduced for attributes that are assumed to have a known (typically monotone) preference order for the levels. In addition, constraints can be imposed for attributes for which respondent-specific preference orders for the levels are obtained in the self-explicated data. If these respondent-supplied orders can be assumed to be valid, such constraints give ACA an (additional) edge over full profiles that can only be matched if the full profile method is expanded to include (some) self-explicated data.

**Table 1**

**Characteristics of ACA and Full Profiles**

<b><u>ACA</u></b>	<b><u>Full Profile</u></b>
- computer-interactive	- any desired form of data collection
- combines design, data collection, analysis, and market simulation	- design, collection, analysis, and market simulation are often separated
- can accommodate a large number of attributes	- restricted to about six attributes, unless bridging designs are used
- objects never fully specified (two to five attributes)	- objects specified on all attributes
- combines self-explicated data with paired comparison intensity ratings	- fully decompositional approach
- stimulus design not orthogonal (statistically inefficient)	- stimulus design typically orthogonal
- paired comparisons are “close” in overall utility	- orthogonal arrays produce profiles that may differ greatly from each other in overall utility
- paired comparisons adapted to respondent specific prior evaluations	- usually same set of full profiles for all respondents
- could impose a <i>priori</i> constraints for all respondents, and/or respondent-specific constraints using self-explicated data, on the estimated partworths	- a <i>priori</i> constraints could be imposed on the estimated partworths
- expected to be subject to attribute level effects (but to a smaller degree than full profile)	- subject to systematic effects, e.g., results depend on the order of attributes, and the number of attribute levels

It is also of interest to point out that the full profile method is subject to some systematic design effects. For example, Johnson (1989) indicates that the relative importance of attributes (derived from the partworths) is influenced by the order in which the attributes appear in the full profiles. That is, if the same attribute appears first for every respondent, it tends to receive more attention than if it appears later. Such a systematic effect should not exist in ACA, because the paired comparison profiles are not presented based on a constant or systematic order of the attributes. Also, Wittink *et al.* (1989) mention that full profiles (as well as tradeoff matrices) are subject to systematic effects due to the number of levels on which attributes are defined. Attribute level effects should be reduced for ACA, because of the strong influence of the self-explicated data on the final partworths (Green *et al.* 1991) and the use of preference intensity judgments.

In ACA, the self-explicated data (preference orders for attribute levels and importance ratings for the differences between best and worst levels of attributes) are used to obtain an initial set of partworths. These initial partworths are defined such that the difference between the values for the best and worst levels of an attribute corresponds to the stated importance of the attribute. Intermediate attribute levels obtain values that correspond to the preference order for the levels. For example, if the attribute warranty has three levels and its importance equals 2 on a 4-point scale, the initial partworths are +1, 0 and -1, respectively for the best, intermediate and worst levels.

One may argue that the paired comparison intensity judgments are or should be used 1) to differentially stretch the attributes' importances, and 2) to modify the assumption of equal utility distances between successive attribute levels. However, there is nothing in the estimation procedure that prevents the updated partworths from altering the assumed or specified preference order for the levels. Especially for *a priori* known (e.g., monotone) preference orders for the levels of an attribute, it is very likely that imposing order constraints on the final partworths will improve the results. Green *et al.* (1991) in their discussion of ACA focus primarily on the extent to which the paired comparison intensity values are commensurate with the initial (or previously updated) partworths. They suggest that ACA results can be improved by differentially weighting the self-explicated data and the paired comparison preferences. Johnson (1991) notes that the nature of differential weights for optimal improvement in predictive validity varies between applications. Thus, improvements in ACA performance are possible if the differential weights can be found (optimized) for each application separately.

## **Experimental Design**

The primary objective of this study is to obtain an empirical comparison of the ACA and full profile conjoint methods. In particular, we are interested in determining the predictive validity of the two methods as determined for four choice (validation) tasks.

We chose refrigerators as the product category. (We could imagine that the absolute and relative predictive validities of the two conjoint methods depend on the product category. However, it is not at all clear in what manner or for what reasons the possible superiority of one method over another depends on the product category.) To enhance the generalizability of the results we varied the number of attributes. (As the number of attributes increases the full profile method becomes more difficult to apply.) Specifically, half the respondents saw five, while the other half saw nine. All nine attributes are described in Table 2.

**Table 2**

**Refrigerator Attributes**

- A. Brand Name - General Electric, Sears/Kenmore, Whirlpool
- B. Capacity - cubic feet: 19, 20\*, 21\*, 22
- C. Energy Cost - annual: \$70, \$80\*, \$90\*, \$100
- D. Compressor - extremely quiet, somewhat quiet\*, somewhat noisy\*, extremely noisy
- E. Price - \$700, \$850\*, \$1,000\*, \$1,150
- F. Design - freezer on left (side by side), freezer on top
- G. Warranty - 1 year, 3 years
- H. Refrigerant - soft CFC (environmentally safe), chlorofluoro-hydrocarbon (hurts environment)
- I. Dispenser - dispenses ice and water through the door, no door dispenser for ice or water

In full profile, half the respondents saw all four levels for attributes B (Capacity) and E (Price), and only the extreme levels for attributes C (Energy Cost) and D (Compressor). The other half saw four levels for C and D, but two for B and E. In ACA, attributes D and E were each given four levels for half the respondents, and attributes B and C had four levels for the other half.

Each respondent provided both full profile ratings and ACA judgments. However, because the predictive validity of the (holdout) choices for each method may depend on the order in which the methods were applied, we used both orders. Thus, half the respondents completed the full profile task first, while the other half did the ACA task first. Both tasks were administered by computer. This manipulation allows us to determine the extent to which task order effects exist and to adjust the results for their presence.

A third manipulation consisted of the number of attribute levels used for four of the attributes. These four attributes—capacity, energy cost, compressor noise and price—were included in all cells of the experimental design. For example, in the full profile design, half the respondents were exposed to four capacity, four price, two energy cost and two compressor levels. The other half saw two capacity, two price, four energy cost and four compressor levels. For all four attributes, respondents' preferences for the levels were assumed to be monotone. Extreme levels were held constant in this manipulation. The primary reason for this manipulation in the study is the expected effect on derived attribute importances (Wittink 1990). We do not address results on the issue of attribute level effects in this paper.


The fourth manipulation involved the order in which the attributes were listed in the full profile method. Attribute order effects have been reported by Johnson (1989). Attributes may receive more attention from respondents if they are listed first (or last) than if they occur closer to the middle of the list. A direct way to investigate this would be to shuffle the attribute order. However, some attributes may have a “natural” place in the order. We decided that all conjoint tasks should have brand name as the first and price as the last attribute. Thus, we manipulated only the order of the remaining attributes. We note that this constraint reduces the opportunity for systematic order effects to occur (Our reason for this constraint is that in this study we were not interested in learning the maximum magnitude of such an attribute order effect in the full profile method. Rather, we wanted to maintain realism in the task characteristics while allowing for an effect to exist. Of course, a validation task that has exactly the same attribute order as the conjoint task would show inflated validities if 1) there is an attribute order effect, and 2) the order in which attributes are attended to in the marketplace differs from the order used in the conjoint method). No systematic effects of this manipulation on predictive validity were detected. An overview of the experimental design is shown in Figure 1.

No attribute order manipulation is indicated for the ACA method. The reason for this is that the order (and identity) of attributes varies between the pairs of profiles (and between respondents). That is, there is no fixed attribute order for the paired profile intensity ratings. The order in which attributes are evaluated in the self-explicated part of ACA is the same as it is in the full profiles for half the respondents. However, the order of the attributes for the objects in the validation choice tasks differs from either order used for the full profiles (and hence for the self-explicated part of ACA).


**Figure 1**

**Experimental Design**<sup>1</sup>

I. Order of tasks administered

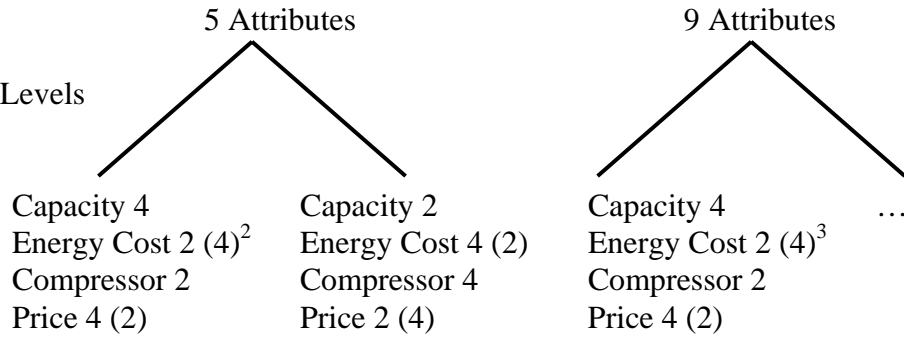


II. Number of Attributes




Price 4 (2)

III. Attribute Levels



IV. Attribute order for Full Profiles

(parentheses indicate  
5 attribute design)



<sup>1</sup> The design is a  $2^4$  design with 16 cells

<sup>2</sup> The number in parentheses is the number of levels used in ACA.

<sup>3</sup> The number in parentheses is the number of levels used in ACA.



Respondent Selection. To obtain broad geographic representation, 400 respondents were selected from a total of 11 cities: Baltimore, Charlotte, Cincinnati, Cleveland, Colorado Springs, Denver, Detroit, Los Angeles, Milwaukee, Philadelphia and Washington, D.C. Each city contributed 36 or 37 respondents. Respondents were selected by intercepting shoppers in super-regional malls. Given four manipulations with two options each, every one of the sixteen cells in the experimental design has 25 respondents. The task characteristics faced by respondents were varied systematically according to the experimental design. This minimized the chance of having systematic differences in respondent characteristics between experimental cells.

To participate in the study, respondents had to have a refrigerator in their home. They were recruited to central locations and were paid \$5 each as compensation for their time. All questions, including those for the full profile task, were asked via computer. Prior to the conjoint tasks, respondents provided information about the brand of refrigerator they owned as well as its age. They were then told to imagine they were shopping for a new refrigerator.

For the full profiles, we used a buying likelihood question: “How likely would you be to buy this refrigerator if you were to buy a new refrigerator today?” A nine-point scale was used (1 = 10% or less (not at all likely); 2 = 20%; ... ; 9 = 90% or more (extremely likely)).

The ACA task followed the standard format of 1) eliciting preferences for the levels of each attribute, if necessary, 2) obtaining importances for the difference between best and worst levels for the attributes, and 3) preference intensity ratings for one of two refrigerators in 12 pairs for designs with 9 attributes, and 6 pairs for designs with 5 attributes. The preference ratings were obtained by asking: “Which refrigerator do you prefer?” A nine-point scale was used varying from 1 = strongly prefer left, to 9 = strongly prefer right.

Validation Tasks. The validation tasks consisted of the following: For each of two pairs of refrigerators, the respondent was asked to indicate which one he or she would be most likely to buy. And, for each of two triples of refrigerators, the respondent indicated which refrigerator he or she would be most likely and (of the remaining two) least likely to buy. Each respondent provided validation data twice (on the identical set of pairs and triples). The validation tasks were separated by the second conjoint task. One important reason for collecting two sets of validation data is that it allows us to assess the reliability or consistency of the choices provided. The predictive validity of either conjoint task cannot exceed the reliability of the choices (Johnson 1989).

In the validation tasks, all refrigerators were defined in terms of the five common attributes. These attributes were listed in the order A, B, E, C, D (see Figure 1). For the attributes with level manipulations, only the two extreme levels (which were included in all cells of the experimental design) were used to define the refrigerators. The refrigerators within each choice set were defined at approximately equal expected average overall utilities, based on our prior judgment.

At the end of the interview session, respondents were asked to answer nine questions about

the two conjoint tasks. The questions asked whether the task “was enjoyable,” “was easy,” “was realistic,” “allowed me to express my opinions,” “took too long to do,” “was frustrating to do,” “asked about too many refrigerators,” “made me feel like just pushing the numbers to get done,” and “had too many features to consider at once.” In addition, respondents provided age and family income information.

## **RESULTS**

For each of 393 respondents (7 respondents provided incomplete data) we estimated two partworth models: one based on the full profile evaluations and another based on the ACA responses. Partworths were obtained based on ordinary least squares. We show results for the holdout choice tasks in Table 3. To provide an overall indication of the predictive validities we converted each triple to three pairs. Thus, with two pairs and two triples twice evaluated we have sixteen pairs per respondent. In this overall sense, full profile has a 68.5 percent and ACA a 72.6 percent hit rate. This difference is statistically significant ( $p < .01$ ).

**Table 3**

**Predictive Validities for Each Method**

		Method		<b><u>p-value</u><sup>1</sup></b>
		<b><u>Full Profile</u></b>	<b><u>ACA</u></b>	
<b><u>Overall</u><sup>2</sup></b>		.69	.73	<.01
<b><u>Task Order</u></b>				
	First	.64	.73	<.01
	Second	.74	.73	-
<b><u>Number of Attributes</u></b>				
	Nine	.66	.71	<.05
	Five	.71	.75	<.05
<b><u>Consistency</u></b>	<b><u>Sample Percent</u></b>			
	≥4	.58	.59	-
	5 or 6	.60	.69	<.01
	7	.70	.74	<.10
	8	.79	.80	-
	100			
<b><u>Choice Task</u></b>				
	Pairs from Triples	.68	.73	<.01
	Pairs	.70	.72	-
<b><u>Triples</u></b>				
	Most likely	.54	.63	<.01
	Least likely	.59	.62	-

---

<sup>1</sup> The p-value is the level of statistical significance for the null hypothesis of no difference between the methods.

<sup>2</sup> For all calculations, except for Triples, the triples were converted into three pairs each. Thus, with choices on two pairs and two triples being provided twice, we have sixteen pairs per respondent.

We have also decomposed the hit rates in several ways. Given that there may be a systematic task order effect, we show the hit rates separately for the estimated partworths obtained from the first and the second conjoint task completed by the respondent. Apparently, the task order has no effect on the ACA hit rates. On the other hand, the full profile hit rates are 64 percent if ACA is completed first and 74 percent if it is the second task. One interpretation of this difference is that the ACA task provides a training opportunity that enhances the consistency of subsequent full profile evaluations. Of course, in practice one would use either full profiles or ACA, and in that context only the results for the first task are relevant. Thus, ACA is substantially better if only one conjoint task is administered to each respondent. It is conceivable, however, that the full profile method improves if a practice or warm-up task is administered first.

The next item in Table 3 is a breakdown by number of attributes. Half the respondents did both conjoint tasks for five attributes, and the other half did both for nine attributes. All choice tasks involved the five common attributes only. The results indicate that the predictive validity is reduced when the conjoint tasks involved four additional attributes. This attenuation is similar for the two conjoint methods. For both five and nine attributes, ACA has higher predictive validities. We note that with nine attributes and sixteen profiles the partworth model is highly saturated. Nevertheless, the increase from five to nine attributes apparently does not reduce the predictive validity more for full profiles than for ACA.

We have also grouped respondents by the degree of consistency in choices. We have choices for each of eight pairs twice, for each respondent. In Table 3, we show predictive validities separately for respondents with low (four or fewer), moderate (five or six), high (seven) and perfect (eight) consistencies. If the consistency is low (unreliable choices), the predictive validities are equally low for the two conjoint methods. Similarly, for respondents with perfect consistency (31 percent of the sample) the predictive validities are equally high. The differences in predictive validity occur for respondents with moderate and high consistencies in choices. One interpretation of this result is that for respondents with moderate to high consistency in their choices (56 percent of the sample) ACA better captures their preferences because of the way in which it simplifies and breaks down the preference elicitation task.

We also note that as the consistency in choices increases, the predictive validities increase for both conjoint tasks. This is in part because 50 percent is the best that can be done by a respondent who is entirely inconsistent. One way to adjust the predictive validities for choice inconsistencies is the following:

For all choices and respondents, the average number of consistent choices equals 6.44 out of 8 (or 80.5 percent). We want to find the hit rates for the two methods when the choices are perfectly consistent. This means adjusting the observed hit rates of .685 for full profile and .726 for ACA by the inconsistent choices for which the hit rates are necessarily .5.

$$\text{Let } (p \text{ con})(adj \text{ hit}_i) + (1 - p \text{ con})(.5) = hit_i$$

where:  $p \text{ con}$  is the proportion of consistent choices (80.5);  
 $adj \text{ hit}_i$  is the hit rate adjusted for inconsistent choices for method  $i$ ;  
 $hit_i$  is the observed hit rate for method  $i$ ;  
 $i = 1, 2$ .

Solving equation (1) for the adjusted hit rates gives .730 for full profiles and .781 for ACA. This adjustment for inconsistent choices increases the predictive validities for both methods and increases the difference slightly.

So far, we have made no distinction between the pairs and triples. There are several reasons why it is useful to examine the triples separately. One, real world choices are not restricted to pairs. Thus, it is of interest to determine the predictive validity of a method when respondents choose, say, the best out of three options. Two, in ACA the paired comparison intensity ratings may be similar to the pairs in the choice tasks.

The separation of the percent hits for all sixteen pairs into twelve pairs constructed from the triples and four pure pairs shows that ACA has higher hit rates for both the pairs derived from the triples (.73) and the pure pairs (.72). However, the difference is more dramatic for the triple-based pairs. Also, the full profile data provide higher hit rates for the pure (.70) than for the triple-based pairs (.68), while the opposite is true for ACA. Overall, these results suggest that ACA has a greater edge over full profiles when the choice task consists of many alternatives compared with a choice task involving two alternatives.

Instead of creating three pairs from each triple, we can also examine the performance of the two methods in a) predicting the most likely to purchase, and b) predicting the least likely to purchase. Here the difference between full profiles and ACA is especially dramatic for the most likely to purchase item (.63 for ACA and .54 for full profiles). For the least likely to purchase item ACA also achieves a higher hit rate, but the difference is not as large (.62 for ACA versus .59 for full profiles). Of course predicting marketplace behavior involves picking the best item, and hence, the results for the most likely item are the best indicator of relative performance for marketplace predictions.

Finally, we show the average respondent perceptions of the two tasks on nine scales in Table 4. On several scales, the average difference between the methods is very close to zero. For example, the responses leaned, on average, toward agreement on the question of task realism (both at 6.29) and toward disagreement on the question whether the task asked about too many refrigerators (3.57 and 3.55) for full profiles and ACA. However, ACA achieved a higher average score on “being enjoyable” ( $p < .05$ ) and a lower average score on “taking too long to do” ( $p < .05$ ). The only other difference approaching statistical significance is that full profile achieved a higher score on “being easy” ( $p < .15$ ).

**Table 4**  
**Respondent Perceptions of the Conjoint Methods**

	<b>Method</b>		<b><u>Difference</u></b>	<b><u>t</u></b>
	<b><u>Full Profile</u></b>	<b><u>ACA</u></b>		
The task ... <sup>1</sup>				
- was enjoyable	5.55	5.81	-0.26	-2.51
- was easy	6.78	6.57	0.21	1.59
- was realistic	6.29	6.29	0.00	0.02
- allowed me to express my opinions	6.43	6.44	-0.01	-0.07
- took too long to do	4.13	3.86	0.27	2.29
- was frustrating to do	3.25	3.31	-0.06	-0.60
- asked about too many refrigerators	3.57	3.55	0.02	0.19
- made me feel just like pushing the numbers to get done	3.47	3.36	0.11	0.95
- had too many features to consider at once	3.77	3.74	0.03	0.20

---

<sup>1</sup> Responses to each question were obtained on a nine-point scale (1 being not at all agreeing and 9 being very much agreeing).

## CONCLUSIONS

We provide evidence on the predictive validity of full profiles and ACA. In our study half the respondents considered questions involving five attributes (four monotone), while the other half had information on nine attributes (seven monotone). The advantage of varying the number of attributes is that it allows us to see if the relative performance of the methods changes with changes in the number of attributes. In either case, however, the number of attributes falls short of the number for which Green and Srinivasan (1990) suggest procedures such as ACA may have an advantage over full profiles.

For both five (when full profile is recommended) and nine attributes, we find that ACA outperforms the full profile method. We have broken down the overall performance results in a number of ways, and observe that ACA maintains an edge each time. The difference is greatest when 1) we use the partworths estimated from the conjoint task performed first by the respondents, 2) we focus on the respondents for whom the holdout choices show moderate to high consistency, 3) we use the triples in the holdout choice tasks, and 4) we consider the predictive validity for the most likely option in the triples. We also have a higher average score for ACA than for full profiles on the task “being enjoyable” and a lower average score on the task “taking too long to do.”

We intend to do a more detailed analysis of the results, for example to document the reliability (statistical significance) of all observed differences. There are also additional experimental manipulations and other effects that still need to be investigated. So far, however, the results quite clearly show that ACA outperforms the full profile method. Of course, the results of this study cannot be assumed to generalize to all contexts for which compositional or decompositional preference measurement is considered. The selection of the product category and the respondents may have a bearing on the results. Also, we administered the full profile method by computer. Seeing and evaluating one profile at a time may make it more difficult for respondents to be consistent in the full profile evaluations. And the finding that the hit rate for full profile is much higher when it is used after ACA suggests that a warmup card sort or attribute importance task may improve the predictive validity of full profile.

## REFERENCES

- Agarwal, Manoj K. and Paul E. Green (1989), "Adaptive Conjoint Analysis Versus Self-Explicated Models: Some Empirical Results," Working paper, State University of New York at Binghamton, December.
- Finkbeiner, Carl T. and Patricia J. Platz (1986) "Computerized Versus Paper and Pencil Methods: A Comparison Study," paper presented at the Association for Consumer Research Conference, Toronto, October.
- Green, Paul E. and V. Srinivasan (1990), "Conjoint Analysis in Marketing Research: New Developments and Directions," *Journal of Marketing*, 54 (October), 3-19.
- Green, Paul E., Abba M. Krieger and Manoj K. Agarwal (1991), "Adaptive Conjoint Analysis: Some Caveats and Suggestions," *Journal of Marketing Research* (in press).
- Johnson, Richard M. (1987), "Adaptive Conjoint Analysis," in: *Sawtooth Software Conference Proceedings*, Ketchum, ID: Sawtooth Software, July, 253-65.
- Johnson, Richard M. (1989), "Assessing the Validity of Conjoint Analysis," *Sawtooth Software Conference Proceedings*, Ketchum, ID: Sawtooth Software, June, 273-80.
- Johnson, Richard M. (1991), "Comment on Green *et al.*," *Journal of Marketing Research* (in press).
- MacBride, J. N. and Richard M. Johnson (1979), "Respondent Reaction to Computer Interactive Interviewing Techniques," paper presented at the ESOMAR Conference.
- Srinivasan, V., Arun K. Jain, and Naresh K. Malhotra (1983), "Improving Predictive Power of Conjoint Analysis by Constrained Parameter Estimation," *Journal of Marketing Research*, 20 (November), 433-8.
- Wittink, Dick R. (1990), "Attribute Level Effects in Conjoint Results: The Problem and Possible Solutions," *First Annual Advanced Research Techniques Forum Proceedings*, Chicago: American Marketing Association, forthcoming.
- Wittink, Dick R., Lakshman Krishnamurthi and David J. Reibstein (1989), "The Effect of Differences in the Number of Attribute Levels on Conjoint Results," *Marketing Letters*, 1, 2, 133-23.