Sawtooth Software

RESEARCH PAPER SERIES

The Benefits of Accounting for Respondent Heterogeneity in Choice Modeling

Bryan K. Orme,
Sawtooth Software, Inc.
1998

530 W. Fir St.
Sequim, WA 98382
(360) 681-2300
www.sawtoothsoftware.com

The Benefits of Accounting for Respondent Heterogeneity in Choice Modeling

Bryan K. Orme

Marketers have long recognized that people are unique. With respect to market research data sets, *heterogeneity* refers to unique preferences or characteristics between segments or from person to person.

If you attended our 1997 Conference, you heard a number of talks addressing the benefits of capturing heterogeneity in CBC analysis, such as more accurate simulation models and better understanding of market structure. You also heard the term IIA (Independence of Irrelevant Alternatives) often referred to as the "Red Bus/Blue Bus Problem." This article will demonstrate that IIA is much more problematic with aggregate logit and much less a problem with methods that recognize respondent heterogeneity, such as Latent Class and ICE (Individual Choice Estimation).

The "Red Bus/Blue Bus" Problem

In the classic "Red Bus/Blue Bus" example, different modes of transportation are available. Assume the following aggregate level utilities for a given set of individuals:

Aggregate Logit Utilities	
Car	0.61
Bus	-0.28
Bicycle	-0.33
Red	0.01
Blue	-0.01

Respondents on average prefer cars to buses or bicycles, and the color of the vehicle is virtually unimportant. Simulated shares of choice for three different modes of transportation are given below:

Aggregate Logit Simulations: Scenario 1	
Car	55.4
Red Bus	23.0
Bicycle	21.6
	100.0

(If you would like to review how shares are calculated using logit utilities, please refer to pages 4-24 and 4-25 of the CBC manual.)

Suppose the bus company wanted to increase traffic on its routes by repainting half of its fleet blue. We would expect that this move shouldn't significantly increase bus ridership, but the aggregate logit simulation suggests differently:

Aggregate Logit Simulations: Scenario 2		
Car	45.1	
Red Bus	18.8	
Blue Bus	18.4	(Net Bus = 37.2%)
Bicycle	17.7	
	100.0	

The net bus ridership has increased from 23.0% to 37.2% (a 62% relative increase), which is not logical. Under aggregate logit simulations, IIA typically results in unrealistic share inflation for very similar (or identical) products.

If we analyze the same data set with Lclass analysis, we get much better results. Lclass detects three segments of individuals with quite different preferences for transportation. The segment sizes and utilities for each group are listed below:

Lclass Utilities

Segment:	1	2	3
Group Size:	16.7%	16.7%	66.6%
Car	-5.9	-1.5	1.5
Bus	2.3	1.4	-0.8
Bicycle	3.6	0.1	-0.8
Red	-0.05	0.07	0.03
Blue	0.05	-0.07	-0.03

Notice that the Lclass utilities are larger in magnitude than the aggregate logit utilities. With logit and Lclass analysis, the better the data fit the responses to choice tasks, the greater the absolute values of the utilities. The Lclass utilities fit the choices much better within each latent class (segment) than a single average set of utilities fits the entire sample.

Let's again simulate results for Scenario 1, this time using the Lclass utilities.

Latent Class Simulations: Scenario 1

Segment	1	2	3	Total
Group Size:	16.7%	16.7%	66.6%	100.0%
Car	0.0	3.9	83.1	56.0
Red Bus	20.6	76.6	8.6	21.9
Bicycle	79.4	19.5	8.3	22.0
	100.0	100.0	100.0	100.0

With Lclass, not only do we get an overall share of choice for the entire sample, but we learn about the composition of the market. Segment 1 favors bicycles; Segment 2 favors buses. The largest segment, Segment 3, favors cars. For this simulation which features dissimilar products, the average (weighted) shares across all segments closely match the shares computed under aggregate logit. With respect to Scenario 1, other than better understanding the composition of the market, Lclass hasn't significantly changed the overall simulation results.

Now let's add blue bus to the simulation.

Latent Class Simulations: Scenario 2

Segment	1	2	3	Total	
Group Size:	16.7%	16.7%	66.6%	100.0%	
Car	0.0	2.4	76.9	51.6	
Red Bus	16.8	45.9	7.9	15.8	
Blue Bus	18.5	40.0	7.5	14.7	(Net Bus = 30.5%)
Bicycle	64.7	11.7	7.7	17.9	
	100.0	100.0	100.0	100.0	

Net share to buses has again increased, but instead of a 62% relative share increase as with aggregate logit, under Lclass the net share for buses increases from 21.9% to 30.5%, representing a 39% increase. Accounting for heterogeneity has reduced the Red Bus/Blue Bus problem.

Where's the magic? Most of the preference for buses comes from Segment 2. In Scenario 1, 76.6% of the share from Segment 2 was cast to the Red Bus. When we add a Blue Bus, there is really very little room for share inflation within that segment since buses already capture nearly all the share.

Imagine that instead of accounting for heterogeneity by fitting utilities to three (or even ten) separate groups, we fit a set of utilities for *each respondent*. ICE, our newest add-on module for CBC, accomplishes just that. The result is a model that is even better than

Lclass at handling the classic Red Bus/Blue Bus simulation problem (and generally superior in terms of holdout predictability). In a real data set we recently collected, the relative share inflation for duplicated products in logit simulations was 61%, 49% and 23% under aggregate logit, Lclass and ICE, respectively.

Even though careful marketing research analysts would never run simulations as unrealistic and methodologically flawed as the Red Bus/Blue Bus situation, most CBC simulations usually involve at least *some* degree of differential similarity among products. These situations can benefit from recognizing heterogeneity using Lclass and ICE.

Cross-Elasticities

Another weakness of aggregate-level logit is its inability to account for cross-elasticities with the standard main-effects or main-effects-plus-interactions models available in CBC. (One can model cross-elasticities with aggregate-level logit, but this requires modeling expertise and software other than CBC.)

Cross-elasticity is defined as the relative percent change in quantity demanded of brand A resulting from a percent change in price of brand B. With aggregate-level logit, when a brand lowers its price, it steals share from other brands in proportion to the other brands' shares. In other words, the cross-elasticities are held constant. If we account for respondent heterogeneity (with Lclass or ICE), some degree of cross-elasticity can be captured and modeled.

Assume two Lclass segments with the following utilities:

Lclass Utilities		
Segment:	1	2
Group Size:	50.0%	50.0%
Cola A	0.5	-0.5
Cola B	0.5	-0.5
Diet Cola C	-0.5	0.5
Diet Cola D	-0.5	0.5
Low Price	1.0	1.0
Medium Price	0.0	0.0
High Price	-1.0	-1.0

Cola A and Cola B are preferred (and highly substitutable) by Segment 1, while diet colas C and D are preferred (and highly substitutable) by Segment 2. Such situations often occur in the real world when competing brands are similarly positioned and appeal to unique segments. To simplify our example, we've made the segments equally sensitive to price.

Scenario 1 shows the simulated shares at the average price:

Scenario 1: All Brands at Medium Price			
Segment:	1	2	Total
Group Size:	50.0%	50.0%	100.0%
Cola A, Medium Price	36.6	13.4	25.0
Cola B, Medium Price	36.6	13.4	25.0
Diet Cola C, Medium Price	13.4	36.6	25.0
Diet Cola D, Medium Price	13.4	36.6	25.0
	100.0	100.0	100.0

In Scenario 2, Cola A lowers its price:

Scenario 2: Cola A Lowers its Price			
Segment:	1	2	Total
Group Size:	50.00%	50.00%	100.00%
Cola A, Low Price	61.0	29.7	45.3
Cola B, Medium Price	22.4	10.9	16.7
Diet Cola C, Medium Price	8.3	29.7	19.0
Diet Cola D, Medium Price	8.3	29.7	19.0
	100.0	100.0	100.0

When Cola A lowers its price, the diet colas each lose (25% - 19% = 6%) share points. Cola B, however, loses a larger amount (25% - 16.7% = 8.3%) share points. Capturing respondent differences in this example has resulted in detecting differential cross-elasticities. Without accounting for the differences between these two segments, simulations would not reveal differential substitutability between the brands.

Interactions

One of the benefits of aggregate-level logit is the ability to model interactions, such as interactions between brand and price. However, if interactions result from differences in preferences between segments, these can also be captured by recognizing heterogeneity with Lclass or ICE without having to include interaction terms. We won't take the space in this article to give a numeric example. Rather, assume a market with two styles of women's dress-pants: a straight leg and a wide leg. Further assume that the straight leg pant tends to appeal to the more price-sensitive segment, and the wide leg appeals to the less price-sensitive segment. If we recognize market heterogeneity when computing utilities, sensitivity simulations can naturally account for an interaction between price and pant styles using only main effects utilities (no interaction terms). Share for the wide leg pants should be less sensitive to price changes than the straight leg.

Summary and Conclusion

Aggregate-level logit has been faulted for its IIA properties. Specifically, aggregate logit modeling can fail when products with differing degrees of similarity are included in simulations. Corrections for product similarity (such as Sawtooth Software's Model 3) can help in such situations, but it is best to begin with an underlying model that is less susceptible to the "Red Bus/Blue Bus" problem. Using Lclass can reduce the need to correct for product similarities under Model 3. For a recent data set we've studied, developing individual-level utilities using ICE largely accounted and corrected for product similarities in logit simulations.

Aggregate logit also cannot account for differential cross-elasticities without customized modeling beyond the standard capabilities of Sawtooth Software's CBC System. Lclass or ICE can account for differential substitutability if such relationships are due to differences among underlying segments or individuals.

CBC has been praised for its ability to detect interactions by pooling data across respondents. To the degree that interactions can be accounted for by differences among segments or individuals, models that recognize heterogeneity can reflect interactions with main-effects only models. If capturing interactions and modeling heterogeneity are *both* of concern, our Lclass Module (which can directly model interactions) can be an excellent approach.