Sawtooth Software

RESEARCH PAPER SERIES

External Effect Adjustments in Conjoint Analysis

Bryan Orme, Sawtooth Software
Rich Johnson, Sawtooth Software

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

External Effect Adjustments in Conjoint Analysis

Bryan Orme and Rich Johnson, Sawtooth Software

March 2006

Background

The market simulator is generally the most practical and powerful deliverable of a conjoint analysis study. It lets the analyst or manager conduct an infinite number of "what-if" scenarios within the context of specific competitive products. The results are expressed in terms of shares of preference summing to 100%. These shares are probabilities, bounded by 0 and 1.0, and interpretable on a ratio scale.

Although many academics refer to simulated shares of preference as "market shares," practitioners usually avoid that label. It is not that those academics somehow believe that conjoint simulator shares are more accurate in predicting real world behavior than practitioners; they just don't have to make client presentations as often. Practitioners face the challenge of communicating the benefits of conjoint simulators without raising unrealistic expectations regarding the tool's ability to forecast sales volume and actual market shares. After all, there are often substantial differences between simulated shares of preference and actual market shares. Conjoint market simulators make a variety of assumptions, including:

- Equal distribution
- Equal awareness
- Perfect information
- Equal time on the market to reach maturity
- Equal effectiveness of sales force and marketing efforts
- That all attributes that influence product choice have been included in the model

Indeed, there are myriad elements that affect actual market shares beyond the fundamental fitness (utility) of the product concept vis-a-vis others on a level playing field. Even when conjoint simulators fail to predict actual market shares, they can be excellent tools for revealing strategic moves for improving market share. But, simulated shares of preference should be taken as *relative* indications of preference rather than indicators of *absolute* sales volume or market share. In the ideal world, managers would recognize the assumptions and accept conjoint simulators for what they do well—ignoring the fact that the shares of preference do not match the actual market shares precisely.

Should We Adjust Shares at All?

We do not live in an ideal world. Many researchers have been able to take the "high road" when presenting simulators and simulation results, while others have faced immense pressure to calibrate the simulated shares to known benchmarks. (Using the word "calibrate" seems to convey a more scientific procedure than simply changing the shares to the desired result through "fudge factors.") Many researchers are paid a premium to apply conjoint data within a larger marketing forecast model, wherein they must account for additional external effects. Conjoint data certainly should be useful in this role.

The purpose of this paper is *not* to encourage or justify the widespread practice of adjusting conjoint simulators in an attempt to account for external effects. Researchers would do well to deliver the simulator as-is, and educate managers regarding the assumptions, proper interpretation, and use of the tool. However, there are cases where researchers must make adjustments to shares to accomplish a specific business purpose. Some types of corrections are perfectly legitimate (such as for distribution and scale factor) and should be done, given the proper data. Other types of corrections are theoretically less justifiable, though perhaps appropriate under certain conditions.

There are few if any published papers on adjustments to conjoint simulators to account for external effects. For many academics, the thought of adjusting conjoint data to match desired base case targets may not only seem heretical, but perhaps also too practical a problem to be of much interest. For practitioners, modifying model parameters is a potentially dangerous topic, but one that needs to be better understood.

Reasons Why Conjoint Simulators Fail to Predict Market Shares

Earlier, we listed some key assumptions of conjoint market simulators. To the degree that these assumptions do not hold in the real world for a specific product category and market, the shares of preference will deviate from actual market shares. There are many other reasons for disconnect between simulated shares and market shares:

Study Design Errors:

- Problems formulating proper attributes and levels
- Choice of conjoint method not compatible with real-world behavior
- Experimental design inadequate to support precise estimates of effects
- Respondents not properly instructed and/or confused by tasks

Data Collection Errors:

- Insufficient sample size
- Improper sampling of people (including modality and non-response bias)
- Time lag between survey data and market share measurement
- Inaccurate measurement of actual market share

Violations of Simulator Assumptions:

- Unequal distribution
- Unequal awareness
- Unequal marketing/sales effectiveness
- Unequal time in market to reach maturity[1]
- Simulator product specifications inconsistent with respondents' perceptions of performance[2]

Respondent Reliability:

- Respondents provide inconsistent answers ("Noise")
- Learning/context effects compromise parameter estimates

Respondent Validity:

- Respondents choose not to answer realistically
- Respondents are unable to answer realistically

Modeling Errors:

- The additive, compensatory assumption in conjoint analysis doesn't accurately capture buyers' actual decision processes
- Unresolved IIA problems/unrecognized heterogeneity
- Failure to account for significant within-concept interaction effects
- Failure to account for significant across-concept interaction effects
- Failure to use the proper market simulation model when converting part worths to shares of preference

Some Approaches to Adjusting Simulated Shares:

Researchers have used various techniques to change the simulated shares of preference to match target base case market shares (in an attempt to account for unexplained variation due to external effects). Some approaches we're aware of are:

Adjustments for awareness

Adjustments for distribution

Scale factor adjustments

A positive multiplicative factor applied to all utilities (also known as the "Exponent" in Sawtooth Software tools)

---

[1] We should note that one can develop factors based on diffusion theory to adjust share for new products. There is quite a lot already available in the literature on this topic. We will not treat it further here, but recognize that such models can also be used in developing external effect factors.

[2] One of the authors (Johnson) has had some experience with this problem. His experience is that when perceptual data are used within simulations, the perceptual data often have greater influence on simulated shares than the part worth utilities. This is not desirable.

Aggregate shares adjustment (Sawtooth Software's Simple Approach)

Positive multiplicative factors applied to aggregate shares, after which shares are re-normalized to sum to 100%

Individual-level adjustments to shares/utilities

Applied as positive multipliers to shares at the individual level, or as individual-level utility adjustments to product concepts

Respondent weighting

A less widely-used procedure that we describe more fully in this article

Adjustments for Awareness

Simulated shares of preference do not match actual market shares due in part to awareness. In the standard conjoint interview, respondents are permitted to choose among all alternatives, even those of which they had previously not been aware. This creates artificial awareness that may not be appropriate to model real world markets.

If we can assume that respondents will not buy brands/products that they are unaware of, then this offers a straightforward mechanism for adjustment. We can simply ask respondents which brands they are aware of and post-process the file of part worth utilities, setting any part worth for brands the respondent is not aware of to an arbitrarily low value (say -15) so that the simulated shares for those brands are zero for this respondent. (This essentially reflects a micro-level correction for product distribution/availability.)

Of course, one could extend this approach to the creation of the conjoint questionnaire itself, customizing the levels to include only the brands of which the respondent is aware. This leads to other opportunities and complexities beyond the scope of this paper.

Lack of awareness may not be such a barrier to purchase in many product categories, where the buyer learns about many options in the natural course of making a purchase. Therefore, using lack of awareness to set a share probability to zero may be too extreme a strategy in many situations.

Adjusting for Unequal Distribution

Unequal distribution is a common reason for simulated shares to deviate from actual market shares. Conjoint interviews typically make all brands and product options available to all respondents. One correction strategy for locally-purchased items is to ask respondents which zip code or other type of region they live in, and then set the part worth utility of any non-available product alternative to some arbitrary low value (such that the alternative will not receive any allocated share in simulations) for each respondent.

When a product has incomplete distribution, one might consider simply reducing its simulated share by some appropriate proportion. But that would assume that all other products gain equally when that product is not available. For example: Assume there are four products (or brands) in a market: A, B, C and D. Further assume that product A receives an overall simulated share of preference of 20% when in competition with B, C, and D for the entire sample. However, A is only available in half the markets. Assuming two equally sized regions, where A is only available in region 1, the proper share is:

Region 1:	A=20% B+C+D=80%
Region 2:	A=0% B+C+D=100%
Overall:	A = 20/(100+100) = 10%

But what about the relative shares for B, C, and D? It might be tempting to deal with this situation just by cutting the share for A in half and redistributing ½ A's share proportionally across B, C and D. However, A did not necessarily compete evenly with B, C and D. Perhaps A was perceived as a close substitute to B, and B would benefit more strongly (than C or D) by A not being available.

Two strategies for this simulation situation are to:

1. Run a simulation with all four products together. Record the shares. Run a separate simulation with only B, C, and D. Record the shares. Average the sets of shares from both runs (where A receives zero share when not available).

2. Replicate the sample. In the second replicate, set the utility of A to a very low value such that no share is allocated to A for any respondent[3]. Run the simulation with competitive set {A, B, C, D} across the total sample (both replicates).

Few modeling situations are as easy as the one described above with just two (or a very few) regions. It is typical to have data like these:

[3] When using Sawtooth Software programs, it is probably most convenient to append a new attribute within the file of respondent part worths (the .HBU file) that includes a new level corresponding to each product to be modified. Use a data processing program script to read in the body of the original .HBU file, modify the data, and write out a new .HBU file. Level one of the new attribute contains an additive utility correction for A, level 2 for B, etc. The additive correction should be very low relative to logit-scaled utilities to force the share to zero for that product for each respondent, such as -15. In the market simulation, specify that A includes level 1 of the appended attribute, B includes level 2, etc. This convention is also useful for more sophisticated adjustments to account for distribution as discussed below. Adjust the "header section" of the .HBU file to account for the additional attribute and levels, and import the .HBU through the Run Manager into a new SMRT project.

Product A available in 90 stores*
Product B available in 120 stores
Product C available in 70 stores
Product D available in 160 stores

**Note: we could as easily say regions or zip codes rather than stores.*

It is tempting just to use the standard approach available in the Sawtooth Software market simulator to adjust the aggregate shares by multiplicative weights proportional to distribution, such as 0.9, 1.2, 0.70, and 1.6 (and re-normalize to 100%). However, these adjustments reflect proportional adjustments to aggregate shares. These products perhaps do not compete equally with one another. And, the patterns of availability of products in the stores (or regions or zip codes) may certainly matter.

If one has more complete data regarding the availability of products within specific regions, this permits the analyst to make adjustments for distribution that also account for unequal substitution effects arising from different combinations of products being available within each region. Assume the following combinations of products available in various stores:

A and D	40 stores
B and D	20 stores
A, B, and D	30 stores
B, C, and D	50 stores
A, B, C, and D	20 stores

Summary:
Product A available in 90 stores
Product B available in 120 stores
Product C available in 70 stores
Product D available in 160 stores

The proportion of stores representing each available combination of products are:

A and D	40/160 = 25.00%
B and D	20/160 = 12.50%
A, B and D	30/160 = 18.75%
B, C and D	50/160 = 31.25%
A, B, C and D	20/160 = 12.50%

Given these data (but no information about which specific respondents shop in which stores), we can simulate respondent shopping trips by assigning respondents to randomly sampled stores offering different combinations of available products. A simple strategy to build such a simulation model starts by randomizing the respondent records in the data file of part worths. Recall above that A and D are the only products available in 25% of the stores. Therefore, for the first 25% of cases, set the utility of brands B and C to arbitrary low values (using a data processing procedure such as was described earlier). Set the utility of A and C to arbitrary low values for the next 12.50% of cases (where only B and D are available), etc. This procedure will result in lower precision for simulated shares because we are discarding some information for many respondents regarding their relative preferences for brands "not available" to them. One way to recover nearly all the benefit of the total sample with regard to precision of estimates is to replicate the sample multiple times (as many times as can be reasonably managed using your software tools), randomizing the order of cases across the replicates, and eliminating brands according to availability as before for different blocks of the expanded sample, proportional to the number of stores offering each unique combination of brand availability. Then, run the simulation across all replicates.

We can generalize the above procedure as a sampling and simulation exercise based on multiple draws per respondent. Within each draw, we simulate the respondent's purchase by random assignment to an available store (or region), where each store reflects the array of true-to-life available brands. If the stores do not capture roughly equal volume, one could easily extend the approach to account for different sales volume per location, by weighting the probability of assignment to a store proportional to the volume accounted for by that location.

From a data processing standpoint, it may be easier to think about each store (or region) as a separate array of zeros and ones, with number of elements equal to total brands in the study (available brand = 1, not available = 0). Replicate the separate files of stores and respondents multiple times, so there are at least as many stores as replicated respondent records. Randomize (independently) the separate files of stores and respondents, and assign the first store to the first respondent, the second store to the second respondent, etc., until all respondent records have been assigned. Adjust the utilities of non-available brands in each case by assigning an arbitrarily low utility of (say, -15). Read the new set of modified part worths into the market simulator. For example, with 800 original respondents, it might be appropriate to create 80,000 replicated records, so that each respondent is simulated for 100 shopping trips among available stores.

Adjustments based on distribution are quite justifiable. We'd suggest doing it prior to any adjustments described below, given you have access to good distribution data.

Scale Factor Adjustments

Consider a set of simulated and target market shares as follows:

	Simulated	Target
A	27%	19%
B	33%	31%
C	40%	50%

At first glance, it would seem that we are predicting brand B very well, but incorrectly predicting for A and C. Upon further inspection, we have correctly predicted C > B > A. Perhaps the errors in prediction are mainly due to *scale factor*.

Imagine that in reality, we had interviewed 500 human respondents and 500 monkeys. The human respondents actually supplied data that perfectly predicted the target market shares. But, the monkeys had answered randomly, with predicted shares of 33.3%, 33.3%, 33.3%. The net result of adding the monkeys to the human respondents is to add random noise, which "flattens" the simulated shares. Experienced conjoint analysts recognize that they can tune the scale factor (exponent) by multiplying all the utilities by a positive constant, prior to applying a market simulation rule such as logit, or Randomized First Choice. (Note that scale factor adjustments have no effect upon the first-choice simulation rule.)

The real world reflects a lot of random behavior, induced by variety-seeking, out-of-stock conditions, imperfect information, and satisficing. It is quite possible that the amount of random error present in the conjoint laboratory does not match (and is lower than) the amount of random variation in real buyers' decisions. Adjusting the scale factor for conjoint data is an appropriate and theoretically justified method for tuning simulated results to more closely fit actual market shares. Scale factor adjustments uniformly tip the shares to be "steeper" or "flatter." They preserve the original rank order of preference. After making appropriate corrections for awareness and distribution, we'd suggest next tuning the scale factor to best fit target market shares. Tuning for scale factor should occur prior to making further adjustments to shares (such as described below) to account for unexplained error.

Procedure for Accounting for External Effects via Aggregate Share Adjustment

If after making adjustments for awareness, distribution, and scale factor the shares still do not match target shares, then the analyst requiring a base case simulation matching target shares will need to make further adjustments. These external effect adjustments are developed once for the base case scenario, and retained in all subsequent simulations. However, they are less defensible, and represent the attempt to somehow account for unexplained differences often without the use of any meaningful explanatory variables. One common approach that makes no pretense of actually explaining the differences involves changing the simulated shares using aggregate share adjustments. This has been offered (with strong cautions in the manual, we might add) in Sawtooth Software's market simulator for many years.

Assume three products in a simulation, with simulated and target shares as follows:

	Simulated	Target
A	0.10	0.15
B	0.30	0.25
C	0.60	0.60

Compute the ratio of target shares to simulated shares:

	Simulated	Target	Ratio
A	0.10	0.15	1.500
B	0.30	0.25	0.833
C	0.60	0.60	1.000

We simply apply the values in the final column above as multiplicative external effect adjustments to the average simulated shares of preference. These simulated shares of preference usually have resulted from shares first computed at the individual level (and then averaged across respondents). After multiplying the average simulated shares by the external effect adjustments and re-normalizing the shares to sum to 100%, the simulated shares exactly match the target shares.

"Reversal Anomaly" with Aggregate Adjustment Method:

In the course of doing technical support for Sawtooth Software's products, we sometimes receive reports of the software seeming to provide strange results. On a few occasions, users have shown us examples where making an isolated change to a product that decreases its utility not only decreases its share, but decreases the share of another (unchanged) competitive product as well. This is counterintuitive, as making one product worse might be expected to result in all other products gaining at least some share.

We'll illustrate the case with an example, based on a real data set (the "TV" data set) that ships with Sawtooth Software products. Consider the following base case, and the shares that result when Sony2 raises its price by $50:

	Base Case	Sony2 +$50	Change
JVC1	9.49	9.71	+0.22
JVC2	18.65	19.96	+1.31
RCA1	9.06	9.94	+0.88
RCA2	28.52	29.84	+1.32
Sony1	14.39	15.54	+1.15
Sony2	19.89	15.00	-4.89

Assume the following external effect factors, to be applied as adjustments to aggregate shares:

JVC1	1.0
JVC2	1.0
RCA1	1.0
RCA2	1.0
Sony1	1.0
Sony2	0.5

After applying external effect factors, the new results for this base case and after the price increase for Sony2 are:

	Base Case	Sony2 +$50	Change
JVC1	**10.53**	**10.50**	**-0.03**
JVC2	20.71	21.58	+0.87
RCA1	10.07	10.74	+0.67
RCA2	31.67	32.26	+0.59
Sony1	15.98	16.80	+0.82
Sony2	11.04	8.11	-2.93

JVC1's share decreases when Sony2 raises its price! This result is quite counterintuitive and can be disconcerting to analysts and managers alike. This anomalous result stems from the fact that the cross-elasticity between JVC1 and Sony2 is quite low (they are not very substitutable). After Sony2 raises its price, the redistribution of its smaller absolute share from Sony2 to JVC1 (than when Sony2 was $50 cheaper) through the external effect adjustment cannot make up for the small gain in JVC1's share due to the increase in Sony2's price. This result is described in more detail for the curious reader in Appendix A.

Procedure for Accounting for External Effects via Individual-Level Utility Adjustment

Suppose that rather than adjusting the aggregate shares, we want to find an adjustment to each brand's part worths, identical for each respondent and unique for each brand, so that the simulated shares become equal to the target shares. A simple iterative procedure for doing this might involve these steps: Observe the differences between target shares and unadjusted simulated shares. Considering either the ratios or differences for those two sets of numbers, add those ratios or differences to each respondent's brand part worth and re-simulate. The discrepancies between simulated and target shares will generally be closer than before, and the same procedure can be applied iteratively to make the discrepancies as small as desired.

This approach focuses on respondents on-the-cusp of choice behavior when shifting shares from one product alternative to another. Respondents with larger scale factor (respondents with higher certainty) are less affected.

We illustrate this procedure with the following example:

	Simulated Shares	Target Shares
A	0.101	0.150
B	0.190	0.250
C	0.061	0.100
D	0.293	0.200
E	0.150	0.100
F	0.204	0.200

Take the ratio of target shares to simulated shares, and zero-center those ratios, by subtracting off the average of the ratios.

Step 1

	Simulated	Target	Ratio	Zero-centered Ratio	
A	0.101	0.150	1.480	0.355	(1.480 – 1.125, etc.)
B	0.190	0.250	1.314	0.188	
C	0.061	0.100	1.628	0.503	
D	0.293	0.200	0.683	-0.442	
E	0.150	0.100	0.667	-0.458	
F	0.204	0.200	0.980	-0.145	
		Average:	1.125		

The zero-centered ratios[4] are added to each product concept and we re-simulate shares of preference given the adjusted product utilities. The new simulated shares (see **Step 2** table below) are closer to the target shares, but we need to make another step in the desired direction. We update the ratio of target to simulated shares, and zero-center those new ratios:

Step 2

	Simulated	Target	Ratio	Zero-centered Ratio
A	0.129	0.150	1.164	0.159
B	0.238	0.250	1.051	0.046
C	0.095	0.100	1.056	0.051
D	0.222	0.200	0.901	-0.104
E	0.117	0.100	0.858	-0.147
F	0.200	0.200	1.001	-0.004
		Average:	1.005	

[4] It really isn't necessary to zero-center the ratios, since a positive constant may be added to all product utilities without changing the simulated shares under the logit model. This convention merely assures that the product utilities do not become so large that it affects the ability of the computer to deal with extremely large numbers that could result after exponentiation if many iterations are performed. Zero-centering also assures uniqueness of the solution (factoring out the arbitrary constant).

We add the new zero-centered ratio to each product concept's previous utility, and re-simulate using the updated utilities. Therefore, the new utility adjustment for A after Step 2 for each individual is 0.355 + 0.159 = 0.514, etc.

The target and simulated shares after six steps of this procedure are:

	Target	Step0	Step1	Step2	Step3	Step4	Step5	Step6
A	0.150	0.101	0.129	0.141	0.146	0.148	0.149	0.150
B	0.250	0.190	0.238	0.247	0.249	0.249	0.250	0.250
C	0.100	0.061	0.095	0.099	0.100	0.100	0.100	0.100
D	0.200	0.293	0.222	0.206	0.201	0.200	0.200	0.200
E	0.100	0.150	0.117	0.106	0.102	0.101	0.100	0.100
F	0.200	0.204	0.200	0.201	0.201	0.201	0.201	0.200

After just six steps, the target shares are obtained to three decimal places of precision. Please note that this example is based on HB utilities at the individual level, which have much greater scale (larger range in utilities) than for latent class, logit, or ACA utilities. When using utilities with smaller scale factor, you may need to reduce the "step size" by multiplying the zero-centered ratios by a value such as 0.50 or 0.10.

This simple procedure may not be the most efficient method for reaching target shares in the fewest steps, but it seems to work pretty quickly. And, assuming you are not using too large a step size, will achieve the desired results. If using a logit-based simulation method, it will always achieve the desired results to as much precision as you desire, and the utility-adjustment solution is unique. In other words, there is only one such set of globally applied zero-centered utility adjustments that will result in the target simulated shares[5].

Procedure for Accounting for External Effects via Respondent Weighting

It is second nature to market researchers to weight respondents based on demographics to match known population characteristics. A related idea might be applied in conjoint simulators.

Sometimes differences between simulated shares and actual market shares may be due to the particular sample of respondents. For example, the sample may contain too few users of a particular product, or it may contain buyers who shop at outlets that don't carry that product. Problems with sample may be due to the sample frame itself, or due to non-response bias. In such cases, it may be possible to improve simulated results by weighting respondents.

[5] Sawtooth Software users can implement adjustments to product utilities in the market simulator using the same method of appending a new attribute (reflecting the utility adjustment, one level for each product to be adjusted) to the file of part worth utilities as was discussed earlier when dealing with adjustments for distribution.

Rather than directly changing simulated shares or modifying model parameters, we seek a set of respondent weights that will produce simulated shares identical to the target shares. Of course, with only a few products to fit, and a very large number of weights that can be adjusted, there would be an infinite number of different solutions to the problem. We choose among those solutions by attempting to find the solution in which the respondent weights are as close to unity as possible.

We start with respondent weights of unity, and calculate a small change for each respondent in the direction that will minimize the sum of squared differences between simulated and target shares. (In optimization terminology, this is the "gradient" vector.) We continue this process through many iterations. In the end, we have computed a set of weights that make simulated shares match target shares exactly, and have done so in many small steps, each of which produced the largest improvement in matching of aggregate shares with the smallest change in the respondent weights. It is not guaranteed that the respondent weights will all remain acceptable—for example, some weights may become negative if the initial discrepancies between individual and target shares are too large. But if these weights are set to zero, the remaining weights will retain their least squares properties. The actual steps in the iterative approach are as follows:

Consider three products with the following simulated and target shares of preference for a sample of respondents:

	Simulated	Target
A	0.25	0.20
B	0.45	0.40
C	0.30	0.40

The simulated shares are based initially on all respondents having a weight of unity.

Consider the first respondent, with individual-level shares of preference (such as from a logit rule simulation model, BTL, or Randomized First Choice) as follows:

A	0.80
B	0.15
C	0.05

1. Subtract simulated shares of preference for the sample from target shares, obtaining the vector [-0.05, -0.05, 0.10]. This characterizes the direction and magnitude in which we would like to adjust simulated shares to match target shares.

2. For each respondent, subtract the target shares from the simulated shares. For the example respondent above, this results in the vector [0.60, -0.25, -0.35]. This vector describes how this respondent's shares compare to the target shares.

3. For each respondent, multiply the vectors from steps 1 and 2. For this example respondent, this results in the vector [-0.03, 0.0125, -0.035].

4. For each respondent, sum the elements within the vector obtained in step 3. For this example respondent, the result is -0.03 + 0.0125 + -0.035 = -0.0525.

5. For each respondent, add the value obtained in step 4 to the previous weight for this respondent. If any resulting weights are below zero (or another desired lower limit, such as 0.10), set them to zero (or the desired limit, such as 0.10).

6. Normalize the weights obtained in step 5 to sum to the number of total respondents, by multiplying each weight by the number of respondents divided by the sum total of weights across respondents from step 5. These are the new weights for respondents.

7. Update the simulated market shares for the sample, based on the new respondent weights.

8. Repeat steps 1 through 7 until weighted simulation results match target shares to the desired degree of precision.

These steps can be executed relatively easily in a spreadsheet program[6], and a macro can be written which copies the new weights over the old weights at the completion of each iteration (paste special + values). It often takes 100 or more iterations to match target shares to a high degree of precision, so it's imperative to automate the procedure using a macro.

In some situations, you may wish to constrain some products to match target shares, while not constraining other products. Simply zero-out any elements in the vector described in step 3 corresponding to non-constrained products.

There are many sets of weights that can be obtained for which the simulated shares match target shares. However, this iterative procedure finds the one unique solution where the weighted simulation returns target shares while also minimizing the standard deviation of weights across the sample (unless weights are constrained, as in step 5, in which case the standard deviation of weights is not minimized). If the target shares deviate considerably from simulated shares or if there is considerable homogeneity in the sample, no solution using positive weights may be possible that lead to a weighted simulation that matches target shares.

[6] "Solver-Like" Excel plugins may also be used to find respondent weights such that a cell containing the mean squared error between simulated and target shares is minimized. However, as of the time of writing, the standard Excel "Solver" plugin was limited to 200 variables (in this case, respondent weights). So, data sets with more than 200 respondents would require more capacity (more capable programs are readily available on the market). Excel's Solver program produces the same results as the iterative procedure described above.

Often, adjusting simulated shares to exactly match target shares may result in quite extreme respondent weights. One can find a compromise between the variance of the weights and the fit to target shares by stopping after fewer iterations[7].

A Series of Tests:

In addition to adjustments for distribution and awareness, we've described three methods that have been used for adjusting simulated shares of preference to match target market shares. The remainder of this paper is dedicated to investigating how these methods perform in a series of tests. Following the tests, we will summarize with conclusions and recommendations.

Test #1: Reduction of Share for One Product: Substitution Effects

Consider the situation in which we want to reduce one product's share dramatically using external effects. We'll consider a base case scenario with six products:

	Base Share
JVC1	9.49
JVC2	18.65
RCA1	9.06
RCA2	28.52
Sony1	14.39
Sony2	19.89

We have constructed these products such that RCA2 and JVC2 are identical in all ways, except for brand. Thus, they should be highly substitutable. What happens if RCA2 raises its price by $50?

Effect of Price Increase for RCA2

	Base Share	RCA2 +$50	% Gain or Loss	
JVC1	9.49	10.97	+16%	
JVC2	18.65	23.58	+26%	(close substitute for RCA2)
RCA1	9.06	10.68	+18%	
RCA2	**28.52**	**17.11**	**-40%**	
Sony1	14.39	16.44	+14%	
Sony2	19.89	21.22	+7%	(not close substitute for RCA2)

---

[7] A twist on this approach that significantly reduces the ratio of maximum to minimum weights (and avoids negative weights) is to assign respondents into groups according to the product with the highest utility (first choice rule). With a six-product simulation, we find only six weights for six respondent groups such that the simulated shares (under the logit, BTL, or Randomized First Choice) match target shares. While the range of weights decreases relative to the method of finding a unique weight for each respondent, the variance of the weights is increased. The results for all three tests shown here are nearly identical to the method of customized respondent weights.

As expected, simulated share of preference for RCA2 reduces and all other products gain share. But, since we are using individual-level (HB) utilities, the gains are not proportional. JVC2 gains the most (on a percentage basis) and Sony2 gains the least. Another way of stating this is that JVC2 is the closest substitute to RCA2 and Sony2 is the least substitutable. Those people seen as most likely to choose RCA2 are also quite likely to choose JVC2 and very unlikely to choose Sony2.

What would happen if we reduced RCA2's simulated share from 28.52 to 17.11 not through making it less desirable by increasing its price, but through various methods of adjusting for external effects? What would we expect to observe? Should the other products absorb RCA2's lost share in a similar pattern to the substitution effects we observed above? We'll test three different methods for adjusting for external effects, and demonstrate that the results are *very* different. We'll label the three methods as follows:

Agg Adj = Aggregate Adjustment (Sawtooth Software method)
Indiv UtilAdj = Individual-Level Utility Adjustment
Resp Wts = Respondent Weights

Change in Product Shares
When RCA2 Cut from 28.52 to 17.11
via Price Increase or External Effects

	Via Price Increase	Agg Adj	Indiv UtilAdj	Resp Wts
JVC1	+16%	+16%	+14%	+18%
JVC2	+26%	+16%	+26%	–2%
RCA1	+18%	+16%	+17%	+14%
RCA2	**–40%**	**–40%**	**–40%**	**–40%**
Sony1	+14%	+16%	+13%	+21%
Sony2	+7%	+16%	+9%	+29%

The three methods of reducing RCA2's shares lead to *very* different results:

- The Aggregate Adjustment redistributes RCA2's share in proportion to the competitors' shares (each competitor gains equally, in proportion to their previous shares).

- The Individual Utility Adjustment redistributes RCA2's share very much in step with the previous substitution patterns that we saw earlier due to increases in RCA2's price.

- The Respondent Weights approach redistributes RCA2's share exactly *opposite* the patterns suggested by substitution effects.

Which is more correct? It depends on your beliefs regarding the need for adjusting RCA2's share to account for external effects.

Individual-Level Utility Adjustment is more appropriate:

- If the adjustment was due to RCA2's lack of full distribution (or awareness), we should expect JVC2 to satisfy those buyers who would have preferred RCA2 but did not find it available (or were not aware of it). JVC2 picks up RCA2's losses at a quicker rate. (However, we have already argued that corrections for distribution and awareness should be accounted for effectively *prior* to adjusting via external effects using one of these techniques.)

- If RCA2's share was too high because of overstated brand utility (equity), sales force effectiveness, or life stage maturity. Its close substitute (JVC2) should pick up share most rapidly when RCA2 is weaker on these fronts.

Respondent Weighting is more appropriate:

- If RCA2's shares are overstated because its performance features are not as desirable as the model predicts. If RCA2's share is overstated, then products (like JVC2) that share similar features might also be overstated. Thus, when RCA2's shares are reduced, products with similar performance aspects like JVC2 should also see a reduction.

- If RCA2's shares are overstated because the sample reflects too many kinds of respondents that gravitate toward RCA2 (and like products). When respondents who tend to like RCA2 are weighted downward, the share for RCA2 (and close substitutes) should also be reduced. Shares for products that appeal to different kinds of respondents (those that are not very substitutable) should increase more rapidly.

Test #2: Reduction of Share for One Product: Own Elasticity

Let's consider again the original base case scenario, with starting shares of preference:

	Base Share
JVC1	9.49
JVC2	18.65
RCA1	9.06
RCA2	28.52
Sony1	14.39
Sony2	19.89

We can easily estimate the price elasticity of demand for the products (%Δ Share/%Δ Price) by increasing the price of each by 10%. The (self) elasticities are:

	Elasticity
JVC1	-2.50
JVC2	-3.82
RCA1	-3.22
RCA2	**-3.21**
Sony1	-3.80
Sony2	-2.09

(Note: even though we are focusing on price elasticity of demand in this example, the same argument holds for sensitivity to other product changes beyond price.)

Rather than reduce RCA2's simulated share through a price increase, let's assume it loses share because it drops a desirable feature. In this case, let's assume it gives up its Picture-in-Picture capability (but continues to charge the same price). Simulated share for RCA2 would drop from 28.52 to 10.32. After this change in product formulation, RCA2's new price elasticity of demand is -4.01, instead of the original -3.21. This is in line with expectations, as a product with lower share and less desirable features should experience an increase in price elasticity.

What would happen if we reduced RCA2's simulated share from 28.52 to 10.32 not through the loss of Picture-in-Picture, but through various methods of external effects? How would the price elasticity be affected? We'll test the three different methods for adjusting for external effects, and again demonstrate that the results are different.

New Price Elasticity for RCA2, When Share Cut from 28.52 to 10.32 via Loss of PIP or External Effect Adjustment

	Loss of PIP	Agg Adj	[8]Indiv UtilAdj	Resp Wts
RCA2	**–4.01%**	**–3.73**	**-3.89**	**-4.29**

All of the methods for adjusting for external effects result in increases in self-elasticity for RCA2 relative to the original -3.21. But, the Respondent Weights method increases the self-elasticity of RCA2 most. Both the Aggregate Adjustment or Individual-Level Utility Adjustment methods increase the price elasticity as well, but at a lesser rate than the loss of PIP.

Which adjustment is more appropriate? Again, it depends on your assumptions regarding why the share should be adjusted, and how the change should be manifest in elasticity.

Aggregate or Individual-Level Utility Adjustments are more appropriate:

[8] Adjusting the share from 28.52 to 10.32 is a dramatic adjustment. The solution that results in the smallest standard deviation of non-negative weights across the sample would have led to 1/5 of the respondents having a zero weight. We constrained the solution such that the smallest respondent weight was 0.10.

- If you expect that changing a product's shares should have a modest impact on the self-elasticity of the product. This is especially the case when corrections principally are due to distribution or awareness. As an example, if a product needs to have its share dramatically reduced because it is only distributed in 1/10th of the market, then its price elasticity should remain constant, despite the fact it has a much lower base. Again, however, adjustments for distribution can be handled in a more appropriate way prior to implementing adjustments for external effects.

Respondent Weights adjustment is more appropriate:

- If you expect that changing a product's share should have a stronger effect on the self-elasticity of the product than is seen with the other methods. If we are missing either some negative or very positive aspects related to the features of the product in our model, then perhaps its price sensitivity should dramatically change as well.

Test #3

In this test, we'll examine the effect of external effects on the relative relationships of self-elasticities and cross-elasticities among products when all product shares are adjusted. Let's consider the previous base case:

	Base Share
JVC1	9.49
JVC2	18.65
RCA1	9.06
RCA2	28.52
Sony1	14.39
Sony2	19.89

Using the market simulator, we can compute the self-elasticities and cross-elasticities resulting from a 10% increase in price for each product.

Elasticities—Effect of Increase in Row Product's Price on Column Product's Share

	JVC1	JVC2	RCA1	RCA2	Sony1	Sony2
JVC1	-2.50	0.35	0.77	0.17	0.27	0.07
JVC2	1.39	-3.82	0.52	1.18	0.77	0.44
RCA1	0.69	0.21	-3.22	0.27	0.42	0.26
RCA2	1.29	2.05	1.47	-3.21	1.18	0.55
Sony1	0.69	0.58	0.95	0.50	-3.80	0.73
Sony2	0.21	0.60	0.81	0.40	0.68	-2.09

The negative values along the diagonal represent self-elasticities. Off-diagonal elements are positive, and reflect cross-elasticities. For example, reading across the first row, an increase of 10% in JVC1's price results in a 3.5% increase in JVC2's share (a cross elasticity of +0.35). The higher the cross-elasticity, the more strongly two products referenced by that cell compete. Also, the larger the share of the row product, the larger the effect of its share changes upon the percentage increase in the column products' shares. Thus, RCA2 (with a base share of nearly 29%) has much higher cross-elasticities across its row than RCA1 (with a base share of around 9%). That is because when RCA2 raises price, it loses larger absolute share (which is redistributed among the other products) than when RCA1 increases price.

As stated before, JVC2 and RCA2 offer the same product features at the same price. When JVC2 increases its price by 10%, RCA2's share increases by 11.8% (a cross-elasticity of +1.18). When RCA2 raises its price by 10%, JVC2's share increases by 20.5% (a cross-elasticity of +2.05).

In this third test, we will apply external effect adjustments to all six products' shares, and compute a new cross-elasticity matrix. We'll compare the new self- and cross-elasticities to the original self- and cross-elasticities prior to adjusting for external effects.

The base case simulated shares and target market share values (hypothetical values used for this example) are as following:

	Base Share	Target Share	Adjustment
JVC1	9.49	14.00	+48%
JVC2	18.65	23.00	+23%
RCA1	9.06	6.00	-34%
RCA2	28.52	22.00	-23%
Sony1	14.39	23.00	+60%
Sony2	19.89	12.00	-40%

After adjustments, the resulting tables of cross elasticities are as follows:

Cross-Elasticities after Aggregate Adjustment

	JVC1	JVC2	RCA1	RCA2	Sony1	Sony2
JVC1	-2.41	0.47	0.90	0.29	0.38	0.19
JVC2	1.58	-3.72	0.70	1.37	0.95	0.61
RCA1	0.54	0.07	-3.31	0.13	0.28	0.12
RCA2	0.89	1.62	1.07	-3.45	0.78	0.17
Sony1	1.13	1.01	1.39	0.93	-3.55	1.17
Sony2	0.00	0.38	0.59	0.18	0.46	-2.25

(Note the near-reversal for Sony2 upon JVC1's share at the bottom-left corner. This is a near-example of the anomaly we discussed earlier.)

Cross-Elasticities after Individual Utility Adjustment

	JVC1	JVC2	RCA1	RCA2	Sony1	Sony2
JVC1	-2.10	0.43	0.94	0.21	0.38	0.09
JVC2	1.30	-3.60	0.69	1.40	0.83	0.55
RCA1	0.29	0.14	-3.14	0.19	0.25	0.19
RCA2	0.78	1.37	1.24	-3.38	0.86	0.41
Sony1	1.05	0.91	1.62	0.92	-3.43	1.19
Sony2	0.10	0.35	0.50	0.24	0.40	-2.16

Cross-Elasticities after Respondent Weights Adjustment[9]

	JVC1	JVC2	RCA1	RCA2	Sony1	Sony2
JVC1	-2.45	0.46	1.62	0.32	0.24	0.12
JVC2	1.47	-3.74	0.82	1.57	0.77	0.68
RCA1	0.58	0.16	-4.25	0.23	0.25	0.25
RCA2	1.02	1.49	1.45	-3.68	0.73	0.57
Sony1	0.76	0.81	1.81	0.80	-3.37	1.63
Sony2	0.13	0.37	0.65	0.29	0.32	-2.33

These tables present a large amount of information that is difficult to grasp unless we dissect the pieces. We'll begin by comparing the self-elasticities in the chart below, before and after the external effect adjustments.

[9] Respondent weights to adjust base case shares to match target shares were as follows: Max weight=2.41, Min Weight=0.28, Standard Deviation=0.77. The ratio of respondent weights from maximum to minimum is 8.6X, which is quite extreme. A clear disadvantage of Respondent Weighting is that the weights may become quite extreme whenever shares need to be adjusted considerably or when the sample is relatively homogeneous.

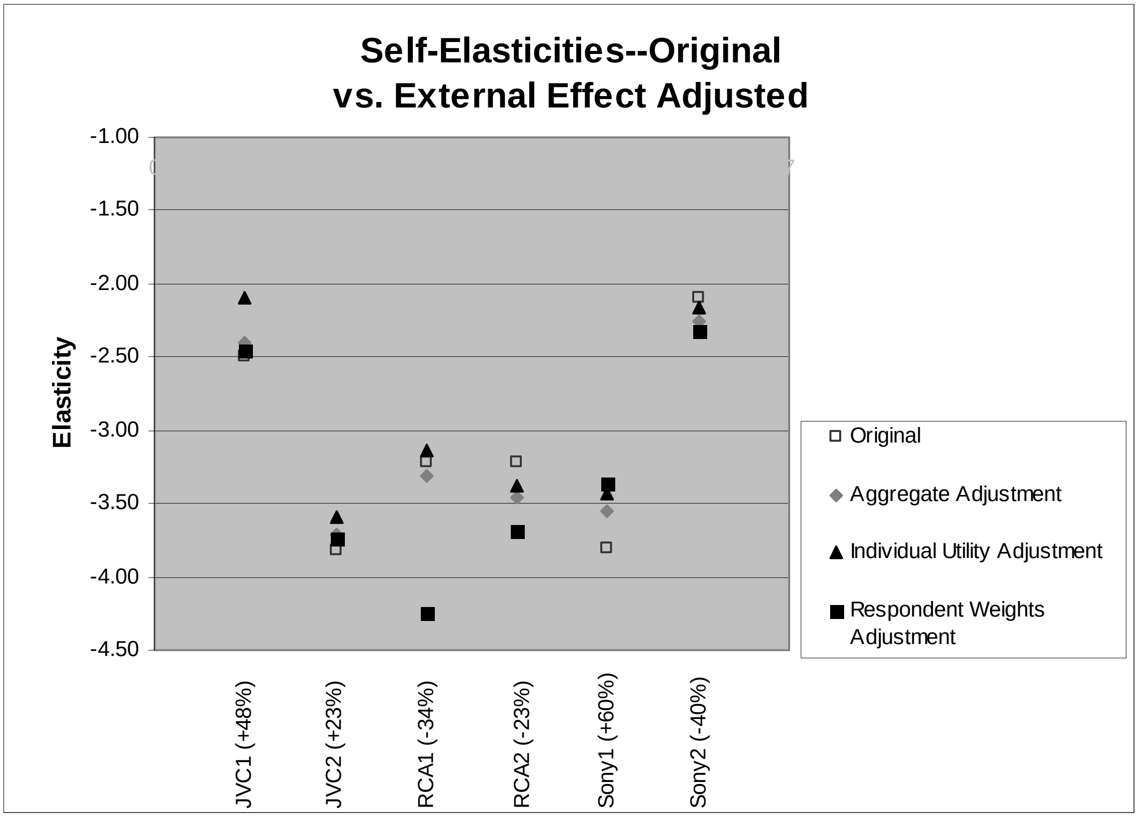

Note that the elasticities have changed very little from the original baseline values when applying Aggregate Adjustment and Individual Utility Adjustments. The Respondent Weights adjustment increases the elasticity quite dramatically for the product whose resulting share after adjustment was smallest (RCA1).

Again, which adjustment method is preferred depends on the reasons why the analyst believes the correction is needed. Adjustments to share due to distribution or awareness should not change elasticities. If the product's share is adjusted because the brand equity or desirability of the features in higher or lower than the model predicts, then changes in share should also result in changes to elasticities. In our example, we have decreased RCA1's share from 9.06 to 6.0 (a 34% decrease in share). The Respondent Weights approach leads to much higher price elasticity for the 6% share RCA1.

Let's now turn our attention to the off-diagonal data entries, the cross-elasticities. The patterns of cross-elasticities are quite similar across all four tables: prior to external effect adjustments, and after three methods for adjusting shares of preference. The correlations among the methods are:

Correlations among Cross-Elasticities

	Prior to Adjustment	Agg Adj	Indiv UtilAdj	Resp Wts
Prior to Adjustment	1.00			
Agg Adj	0.81	1.00		
Indiv UtilAdj	0.78	0.96	1.00	
Resp Wts	0.76	0.90	0.94	1.00

The Aggregate Share adjustment cross-elasticities are most similar to the original cross-elasticities, prior to adjustment. However, we should probably expect that cross-elasticities should change due to the relatively large shifts in shares we imposed upon the products. So, this is not necessarily a mark of success.

The Aggregate and Individual Utility Adjustment methods are highly correlated at 0.96. These approaches indeed result in very similar cross-elasticities after adjustments to shares.

Economic Modeling via Regression Analysis:

It is difficult to judge from the tables of cross-elasticities and from the correlation analysis which external effect adjustment makes most sense. We'll apply a multiple regression model to investigate how consistent the methods are with respect to economic theory. The percentage change in Product j's share due to a change in Product i's price is a function of the shares of products i and j and the amount of similarity between i and j. More formally,

$$\text{Cross\_Elasticity}_{ij} = f\,(\text{Share}_i + \text{Share}_j + \text{Product\_Similarity}_{ij})$$

Here, we define the product similarity of products i and j as the total number of attribute levels shared in common.

Here are the results of the regression equation for the original model (prior to correction with external effects), and after each of the external effect adjustments we've been investigating:

Original Shares (Prior to Adjustment)

R-Squared	0.819		
Constant	-0.079		
	Coeff.	Std Err	T-Value
$Share_i$ (Row Share)	0.040	0.00569	7.11
$Share_j$ (Column Share)	-0.015	0.00569	-2.67
$Similarity_{ij}$ (# Shared Levels)	0.196	0.02902	6.76

Aggregate Share Adjustment

R-Squared	0.800		
Constant	-0.129		
	Coeff.	Std Err	T-Value
$Share_i$ (Row Share)	0.042	0.00674	6.17
$Share_j$ (Column Share)	-0.013	0.00674	-1.91
$Similarity_{ij}$ (# Shared Levels)	0.189	0.03415	5.55

Individual Utility Adjustment

R-Squared	0.827		
Constant	-0.061		
	Coeff.	Std Err	T-Value
$Share_i$ (Row Share)	0.042	0.00599	6.99
$Share_j$ (Column Share)	-0.016	0.00599	-2.73
$Similarity_{ij}$ (# Shared Levels)	0.171	0.03038	5.62

Respondent Weights Adjustment

R-Squared	0.720		
Constant	0.357		
	Coeff.	Std Err	T-Value
$Share_i$ (Row Share)	0.038	0.00889	4.24
$Share_j$ (Column Share)	-0.034	0.00889	-3.83
$Similarity_{ij}$ (# Shared Levels)	0.190	0.04502	4.22

The overall model fit of the best models is quite good (over 80% of the variance in the cross-elasticities explained). In all models, the coefficients have the expected directions and are all significant ($t>1.96$), with one exception, at the 95% confidence level. However, the model fit is noticeably decreased in the case of Respondent Weights adjustments. This analysis would suggest that the other methods more closely align with economic theory—especially the Individual Utility Adjustment method. If the adjustments to product shares (and resulting respondent weights) weren't so extreme, we'd expect the Respondent Weights method to perform better. But to test the comparative qualities of the approaches, it's useful to try more extreme cases.

Below, we have plotted the residuals from the multiple regressions.

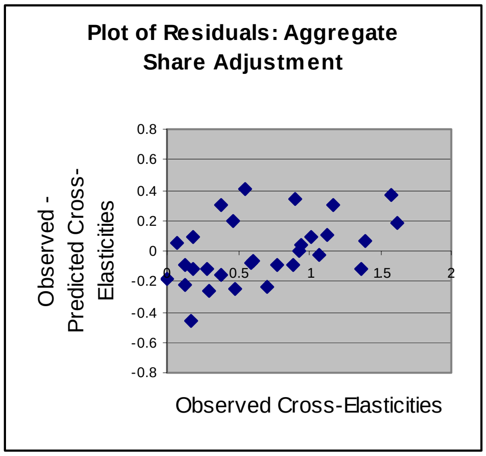

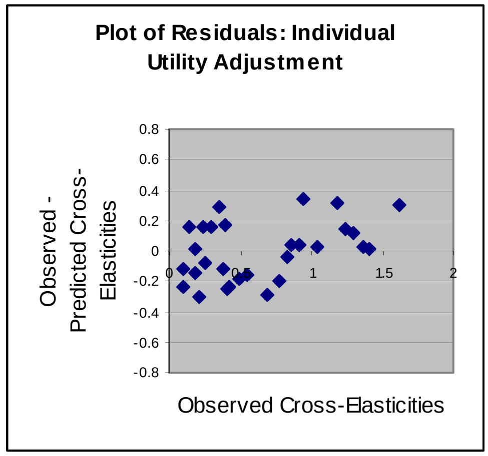

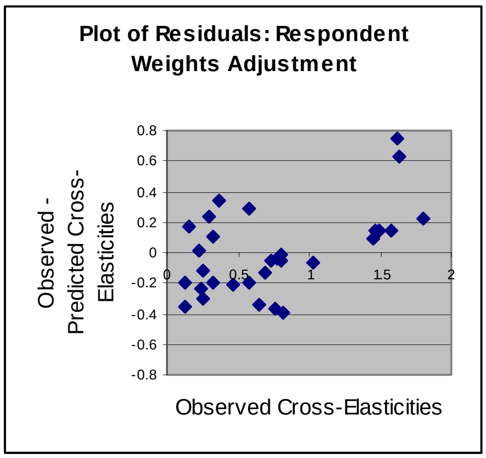

The cases with the largest observed cross-elasticities seem to be slightly overstated, relative to the economic theory regression model. This is especially the case with the Respondent Weights adjustment, with two data points having residuals > 0.6 (associated with the smallest Column share products RCA1@6% and Sony2@12%). The Respondent Weights adjustment dramatically overstates the cross-elasticity in these cases.

Recommendations and Conclusions:

We began by discussing why simulated market shares often do not match target market shares. We covered methods for adjusting for awareness and availability, with special emphasis on a simulation and sampling method for simulating repeated respondent "shopping trips" within sampled "stores/regions," each reflecting available product offerings. We're confident about this approach, and suggest analysts use it whenever the data permit and prior to other external effect adjustments.

Adjustments based on scale factor should always be investigated, given proper market share data, after appropriate corrections have been made for distribution (and possibly awareness). This is a legitimate and often necessary adjustment to control for noise and consistency between two samples and choice environments.

This research suggests that there is no clear winner among the three methods for adjusting shares via external effects. The method you choose certainly depends on the modeling situation and by how much the shares need to be adjusted. But we can characterize some general strengths and weaknesses of the approaches.

- The method of **Aggregate Adjustment** as currently offered in the Sawtooth Software simulator has drawbacks (such as the "reversal anomaly") and is naïve in the way it proportionally re-distributes share. But, it is easy to implement and makes only modest adjustments to self- and cross-elasticities.

- **Individual Level Utility Adjustment** is quite similar in spirit to the Aggregate Adjustment, but leverages the power of individual-level data, focusing changes on respondents on-the-cusp of choice. Shares are not simply re-distributed in proportion to previous shares, as with the Aggregate Adjustment, but are reapportioned according to expected switching patterns among products with varying similarity and correlations in preference. This approach also avoids the "reversal anomaly" that we demonstrated. It also makes only modest adjustments to self- and cross-elasticities, in the expected direction depending on whether the product's share is increased or decreased, which in many cases is the desired outcome.

- **Respondent Weighting** is theoretically best when corrections need to be made due to improper sample composition. However, one should watch for extreme weights, when some respondents are weighted much more heavily (especially 5x or more) than others. In some cases, no solution is possible using positive weights. Respondent re-weighting is not appropriate for corrections dealing with distribution and awareness (which should be handled in other ways). Self-elasticities can become relatively more extreme for adjusted products through respondent weighting. Results from the cross-elasticity test and regression modeling suggest that certain cross-elasticities (substitution effects) may be exaggerated beyond the expectations supported by the economic model we imposed. The variance of the respondent weights directly affects the performance of this adjustment approach. When the variance of respondent weights is not as extreme as in our test case, the results should be greatly improved.

In general, we suggest researchers avoid adjusting simulated shares to match target market shares. But, if you must, we suggest corrections be made in this order:

1. Corrections for availability (do not allow respondents to allocate share to product alternatives that are not available to them).

2. Corrections for awareness, if achieving awareness is a significant hurdle to purchase consideration. Adjustments for awareness are not as straightforward, and we suggestion caution in applying them.

3. Scale Factor. Such adjustments are quite appropriate and defensible (given a good overall conjoint model and that the previous steps are undertaken).

4. Product share adjustments, either through Individual Utility Adjustment or a modest degree of Respondent Weighting, depending on what qualities you believe are lacking in the model and the degree to which shares need adjustment. Hopefully, if you have designed your study well and have properly accounted for distribution, the needed adjustments will be modest.

Appendix A

Aggregate Shares Correction Anomaly

In the body of the paper, we showed how the standard External Effect adjustment as applied in the Sawtooth Software simulator can cause JVC1's share to actually *decrease* when Sony2 raises its price. This example is taken from an HB run for a real data set (the "TV" example data set and accompanying HB run that ships with Sawtooth Software's market simulator). In case the reader is interested to replicate the results using the "TV" data set, the base case scenario is as follows:

Base Case Product Specifications:

	Brand	Screen Size	Sound Quality	Channel Blockout	PIP	Price
JVC Low End	1	1	1	2	2	300
JVC High End	1	2	2	2	2	375
RCA Low End	2	1	2	1	1	325
RCA High End	2	2	2	2	2	375
Sony Low End	3	1	2	2	1	350
Sony High End	3	3	3	1	1	400

We'll repeat the example below:

	Base Case	Sony2 +$50	Change
JVC1	9.49	9.71	+0.22
JVC2	18.65	19.96	+1.31
RCA1	9.06	9.94	+0.88
RCA2	28.52	29.84	+1.32
Sony1	14.39	15.54	+1.15
Sony2	19.89	15.00	-4.89

Assume the following external effects:

JVC1	1.0
JVC2	1.0
RCA1	1.0
RCA2	1.0
Sony1	1.0
Sony2	0.5

After applying external effects, the new results for this base case and price increase for Sony2 are:

	Base Case	Sony2 +$50	Change
JVC1	10.53	10.50	-0.03
JVC2	20.71	21.58	+0.87
RCA1	10.07	10.74	+0.67
RCA2	31.67	32.26	+0.59
Sony1	15.98	16.80	+0.82
Sony2	11.04	8.11	-2.93

JVC1's share decreases when Sony2 raises its price! This result is quite counterintuitive and can be disconcerting to analysts and managers alike.

How does this happen? Let's consider the original simulation, prior to applying external effects. The respondents who prefer Sony2 are not very likely to switch into JVC1 when Sony2 raises its price. Note that JVC1 only gains +0.22 in share with Sony2's price increase (9.71 – 9.49 = 0.22).

Under the external effect adjustment, Sony2's shares are reduced by a multiplicative factor of 0.5. The lost share from Sony2 is then distributed among the other offerings in proportion to their original shares:

	Base Case		Ext. Eff.			Renormalize to 100%	
JVC1	9.49	x	1.0	=	9.49	9.49 / 90.06 * 100 =	10.53
JVC2	18.65	x	1.0	=	18.65	18.65 / 90.06 * 100 =	20.71
RCA1	9.06	x	1.0	=	9.06	9.06 / 90.06 * 100 =	10.07
RCA2	28.52	x	1.0	=	28.52	28.52 / 90.06 * 100 =	31.67
Sony1	14.39	x	1.0	=	14.39	14.39 / 90.06 * 100 =	15.98
Sony2	19.89	x	0.5	=	9.95	9.95 / 90.06 * 100 =	11.04
			Totals:		90.06		100.00

Thus, the loss in share for Sony2 due to the 0.5 external effect adjustment (see last row of the table above) is redistributed according to the proportions of the competitive products, and JVC1 picks up 1.04 share points (from 9.49 to 10.53).

After Sony2 increases its price, its unadjusted share decreases from 19.89 to 15.00. The shares are redistributed according to the External Effect multipliers as below:

	Sony2 + $50		Ext. Eff.			Renormalize to 100%	
JVC1	9.71	x	1.0	=	9.71	9.71 / 92.50 * 100 =	10.50
JVC2	19.96	x	1.0	=	19.96	19.96 / 92.50 * 100 =	21.58
RCA1	9.94	x	1.0	=	9.94	9.94 / 92.50 * 100 =	10.74
RCA2	29.84	x	1.0	=	29.84	29.84 / 92.50 * 100 =	32.26
Sony1	15.54	x	1.0	=	15.54	15.54 / 92.50 * 100 =	16.80
Sony2	15.00	x	0.5	=	7.50	7.50 / 92.50 * 100 =	8.11
			Totals:		92.50		100.00

Prior to the price increase by Sony2, a larger absolute share from Sony2 was shifted to the other brands due to external effect adjustments, resulting in JVC1 picking up 10.53 – 9.49 = 1.04 share points. After Sony2's price increase, Sony2 of course loses overall share. When Sony2's new lower share is shifted to the other brands due to the external effect factors, it results in JVC1 picking up 0.79 share points (10.50 – 9.71), which is 0.25 fewer points than before due to the external effect adjustment. This loss for JVC1 cannot counteract its relatively small gain in share of +0.22 (9.71 – 9.49) due to Sony2's price increase prior to applying external effects. The net change to JVC1 is 0.22 – 0.25 = -0.03.