Client Conjoint Simulator
(CCS)

(Updated 20 October 05)

Sawtooth Software, Inc.
Sequim, WA
http://www.sawtoothsoftware.com

In this manual, we refer to product names that are trademarked. Windows, Windows 95, Windows 98, Windows 2000, Windows XP, Windows NT, Excel, PowerPoint, and Word are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

About Technical Support

We've designed this manual and the on-line help within the Client Conjoint Simulator (CCS) to teach you about how to use the software. If you have questions about issues not covered in this manual, we suggest that you contact the consultant or research supplier that has conducted the conjoint project and delivered the data for your use. Such individual(s) understand your project and business better than we at Sawtooth Software and are in a better position to assist you.

If you would like to contact Sawtooth Software, you may contact us at support@sawtoothsoftware.com or 360/681-2300 (voice) or 360/681-2400 (fax). Outside of the U.S., contact your Sawtooth Software representative for support.

Table of Contents

1 Chapter 1: Installing the Client Conjoint Simulator

1.1 Installation and Removal

Introduction

The Client Conjoint Simulator (CCS) is a software tool for analyzing conjoint/choice data and simulating market choices. It offers a subset of the capabilities of Sawtooth Software's SMRT suite (Suite of Marketing Research Tools). The typical pattern is for a consultant to use the full capabilities of the conjoint/choice software, collect the data and estimate the part-worths (preference utilities), and then provide market simulation capabilities to the client through the Client Conjoint Simulator.

The CCS is licensed to a business entity and involves a licensing fee and agreement. The CCS is offered with a very flexible license. It can be installed on an unlimited number of PCs with an unlimited number of concurrent users within the licensed organization.

You may distribute or print additional copies of this manual (either in electronic or print form).

System Requirements

The CCS runs under Windows 95/98/2000/XP or NT. The minimal system configurations able to run those operating systems are sufficient for running the CCS. Some computational steps within the CCS are quite intensive, so we recommend the fastest processors available. By today's standards a 2 GHz or faster computer is more than adequate.

Installation

Before installing the CCS (SMRT software platform), we suggest that you first terminate any programs that are already running. Installing (or re-installing) the CCS will not delete or overwrite any existing CCS data or study-related files.

1. Insert the CD-ROM installation disk into your CD-ROM drive. (Or click the link to install the system from our website). If faced with a menu with different choices of software to install, choose SMRT.

2. The installation program should automatically start.

3. Follow the instructions. By default, the CCS (SMRT) will be installed to the following directory: **c:\program files\sawtooth software\SMRT**. You can change this directory if you want.

The consultant that prepared the conjoint data for you should have given you five additional files that should be copied onto your hard drive. (For the purpose of illustration, we'll assume a study name of "Study1." A study is made up of a collection of files needed to conduct analysis for a particular project. The study name identifies the name of the project and the prefix for the file names used.)

The list of study-related files to be copied is as follows:

Study1.SMT
Study1.UCS

Study1.QNR
Study1.DAT
Study1.IDX

Running the Client Conjoint Simulator

Click *Start | Programs | Sawtooth Software | Sawtooth Software SMRT*, and the main menu is displayed. To open a study, select *File | Open*, and browse to the directory containing the studyname.SMT file. Double-click the studyname.SMT file.

Removing the Client Conjoint Simulator from Your PC

Windows operating systems do not let you remove software just by deleting a directory. Windows applications leave traces of themselves in many places, and to remove them all you should use the Uninstall procedure. However, if you try to uninstall the CCS (SMRT) while it is still running, the uninstall utility will not be able to do a complete job, and many files will be left on your disk.

IMPORTANT NOTE: Before uninstalling the CCS, make sure it is not currently running.

Uninstalling the Client Conjoint Simulator

After you have exited CCS (SMRT):

1.   Click *Start | Settings | Control Panel*

2.   Double-click *Add/Remove Programs*

3.   Select *Sawtooth Software SMRT*

4.   Click *Add/Remove…*, and follow the instructions.

If you installed study-related files before uninstalling the CCS, the study data and all other related study files will *not* be removed by the uninstall program. This protects you from losing important data or files related to studies you may be working on, or studies you've already completed. If you no longer want these files, you can delete them manually.

2

Chapter 2: Using Simulators to Answer Strategic Questions

2.1 Overview

The Market Simulator is usually considered the most important tool resulting from a conjoint project. The simulator is used to convert raw conjoint (part-worth utility) data into something much more managerially useful: simulated market choices. Products can be introduced within a simulated market scenario and the simulator reports the percent of respondents projected to choose each. A market simulator lets an analyst or manager conduct what-if games to investigate issues such as new product design, product positioning, and pricing strategy.

A Warning about Interpreting the Output of Market Simulators

Under very controlled conditions (such as markets with equal information and distribution), market simulators often report results that closely match long-range equilibrium market shares. However, conjoint part-worth utilities cannot account for many real-world factors that shape market shares, such as length of time on the market, distribution, out-of-stock conditions, advertising, effectiveness of sales force, and awareness. Conjoint analysis predictions also assume that all relevant attributes that influence share have been measured. Therefore, the share of preference predictions usually should not be interpreted as market shares, but as *relative* indications of preference.

Divorcing oneself from the idea that conjoint simulations predict market shares is one of the most important steps to getting value from a conjoint analysis study and the resulting simulator. While "external effect" factors can be built into the simulation model to tune conjoint shares of preference to match market shares, we suggest avoiding this temptation if at all possible. No matter how carefully conjoint predictions are calibrated to the market, the researcher may one day be embarrassed by differences that remain.

2.2   What Is a Market Simulation?

A conjoint study leads to a set of utilities (part-worths) that quantify respondents' preferences for each level of each attribute. These part-worths can be analyzed in a number of ways. You can examine each respondent's part-worths (but this task could become overwhelming). You might summarize the average part-worth utilities, or compute average importances (described in Chapter 6). You could create graphs and charts to display that information, but to many it might seem somewhat abstract and difficult to grasp. Examining average responses could also fail to detect important segments of the market that have unique and targetable preferences.

A good market simulator is like having all of your respondents gathered in one room for the sole purpose of voting on product concepts and competitive scenarios (defined in terms of the attribute levels you measured) you show them. You walk into the room, show them a market scenario (i.e. products A, B and C), and they vote for the one(s) they prefer. Millions of potential products and market situations could be evaluated, and your captive audience would never get tired, ask for lunch breaks, or require you to pay them by the hour.

How does a market simulator work? Let's suppose we had a way (such as through conjoint or choice analysis) to quantify how much people liked the different qualities of ice cream cones. Let's refer to those preferences as part-worth utilities, and assume the following values for a given respondent:

	Utility
Chocolate	0
Vanilla	30
Strawberry	40
$0.60	50
$0.80	25
$1.00	0

Using those utility values, we could predict how he would choose between a vanilla cone for $0.80 or a strawberry cone for $1.00.

    Vanilla (30 utiles) + $0.80 (25 utiles)   =   55 utiles
    Strawberry (40 utiles) + $1.00 (0 utiles) =   40 utiles

We'd predict he would prefer the vanilla cone. If we had data for 500 respondents, we could count the number of times each of the two cones was preferred, and compute a "Share of Preference," also referred to as a "Share of Choice":

	Share of Choice
Vanilla @ $0.80	300/500 = 0.60
Strawberry @ $1.00	200/500 = 0.40

In our hypothetical market simulation, 60% of the respondents preferred the vanilla, and 40% the strawberry cone. This illustrates the most simple simulation approach, referred to as the First Choice model.

2.3   Why Conduct Market Simulations?

Looking only at average preferences (part-worth utilities) can mask important market forces caused by patterns of preference at the segment or individual level. Marketers are often not interested in averages, but in the targetable, idiosyncratic behavior of segments or individuals.

For example, consider the following three respondents, and their preferences (utilities) for color:

	Utilities for Color		
	Blue	**Red**	**Yellow**
Respondent A	50	40	10
Respondent B	0	65	75
Respondent C	<u>40</u>	<u>30</u>	<u>20</u>
Average:	**30**	**45**	**35**

Looking only at average preferences, we would pronounce that red is the most preferred color, followed by yellow. However, if one of each color was offered to each respondent, red would *never* be chosen under the First Choice model, yellow would be chosen once, and blue twice — the exact opposite of what aggregate part-worth utilities suggest. While this is a hypothetical example, it demonstrates that average part-worth utilities do not always tell the whole story. Many similar, complex effects can be discovered only through conducting simulations.

Some reasons for conducting conjoint simulations include:

1. Conjoint simulations transform raw utility data into a managerially useful and appealing model: that of predicting market choice (Share of Preference) for different products. Under the proper conditions, shares of preference quite closely track with the idea of market share — something most every marketer cares about.
2. As demonstrated earlier, conjoint simulations can capture idiosyncratic preferences occurring at the individual or group level. These "hidden" effects can have a significant impact on preference for products in market scenarios. When multiple product offerings have been designed to appeal to unique segments of the market, capturing such effects is especially important for accurately predicting preference.
3. Conjoint simulations can reveal differential substitutability (cannibalism/cross-elasticity effects) between different brands or product features. If two brands are valued highly by the same respondents (have correlated preferences), these brands will tend to compete more closely. Product enhancements by one of these brands will result in more relative share being lost by the correlated brand than by other less similar brands within the same simulation. Examining aggregate part-worth utilities cannot reveal these important relationships.
4. Conjoint simulations can reveal interaction effects between attributes. If the same respondents that strongly prefer the premium brand are also less price sensitive than those who are more likely to gravitate toward a discount brand, sensitivity simulations will reflect a lower price elasticity for the premium relative to the discount brand. A similar interaction effect can occur between many other types of attributes: such as model style and color.

Note when using CBC data: It is important to note that complex effects other than two-way interactions such as cross-effects cannot be reflected using the model of aggregate-level logit offered by our CBC system. Latent Class is a technique for estimating part-worth utilities and reflecting respondent differences at the group/segment level, and CBC/HB (Hierarchical Bayes) and ICE (Individual Choice Estimation) are ways to estimate utilities at the individual level for CBC data. Because they are built on individual-level preferences, simulators based on these models are able to reflect the important and complex behaviors mentioned earlier. It is not surprising that Latent Class and CBC/HB have been shown to outperform similarly-defined aggregate level logit models in terms of predictive validity.

ACA (Adaptive Conjoint Analysis) and CVA (Full-Profile Conjoint Analysis) capture respondent-by-respondent preferences and are thus very useful inputs to this and other market simulation models.

2.4   Typical Questions for Conjoint Simulators

Typical Strategic Questions for Conjoint Simulators

There are many marketing strategies that can be investigated with the conjoint simulator. Here are three of the most common:

1. Given a current competitive environment, what product should I offer to maximize interest in my offering? How can I modify an existing product to capture more relative demand? A Market Simulator lets you input multiple products and place them in simulated competition one with another. Each product is defined using the attribute levels measured in the conjoint study (brands, colors, prices, speeds, warrantees, etc.). Therefore, if you have measured the relevant brands and features offered in the market, you can simulate a realistic market scenario within the Market Simulator. Within that market scenario, you can add a new product and see how well it competes. If the goal is to maximize share, offering the best features at the lowest price is often the trivial solution. The Market Simulator focuses on the demand side of the marketing equation; but it is also important to pay attention to the supply side and take the costs of producing different products/services into consideration. If you have cost information available to you, the Market Simulator permits you to investigate the incremental benefits of different features of a product relative to the cost of offering them. (The CCS does not provide an automatic way to input cost information, but you can divide the incremental gain in share of preference by the incremental cost to produce a profitability index.)

2. What is the relative price sensitivity of different brands? If I raise my price by 10%, how will it affect my brand? How will it affect competitors' brands? You can conduct "sensitivity analysis" for attributes such as price using the Market Simulator to generate relative demand curves. The approach involves holding all other brands at a constant price and changing the price of a single brand, recording the relative share at each point for that brand along the price continuum. If your conjoint data reflect respondent differences (all models offered by Sawtooth Software except for aggregate CBC logit), differential cross-elasticities can be investigated through this approach.

3. What portfolio of products can I offer to appeal to different market segments and maximize overall share? The consultant or research firm that conducted your study and delivered the data to you may have merged segmentation variables (such as demographics or firmographics). If you have segmentation information, you can investigate product formulations that appeal to different groups of respondents. It is likely that by designing products that appeal uniquely to targetable segments that you can increase overall share for your product line or occupy a niche that is not currently being served.

These three strategic questions and the simulation strategies for responding to them are illustrated further in Chapter 5.

3   Chapter 3: Market Simulation Theory and Models

3.1   Base Case Scenario

Introduction

Market Simulations provide a very useful way to use conjoint/choice data. Simulations provide an intuitive tool to move from the esoteric realm of part-worth estimates/effects toward the practical world of predicting buyer behavior for specific market situations. Before using the Client Conjoint Simulator (CCS), you should become familiar with some terminology, issues and theory.

This chapter will discuss setting up a base case scenario, the scaling of simulation results, the Red-Bus/Blue-Bus problem, and the five simulation models offered by the CCS. Finally, External Effects and the use of the None weight are introduced.

Base Case Scenario

Usually the first step in using the market simulator is to define a "Base Case" scenario. A base case typically reflects a current (or future) market scenario: your brand vs. the relevant competition. If there is no relevant competition, or your conjoint study was designed to model only your product, the base case may be a single product, reflecting a likely configuration.

The Market Simulator lets you input the market scenario in a grid format, where the products are rows (you can have up to 100), and the attributes are columns (you can have up to 30).

Product	Attribute 1 (Brand)	Attribute 2 (Package)	Attribute 3 (Color)
Product A	1	2	3
Product B	2	1	2
Product C	3	3	1

You provide text labels in the first column to identify the products. In the attribute cells, you type the numeric value associated with different attributes. For example, Product B is defined by level 2 of Brand, level 1 of Package and level 2 of Color. (The Market Simulator displays the list of attributes and the codes associated with each level as you enter values in the grid.)

After defining the market scenario, you should decide which simulation method is appropriate for your data and the types of strategic questions you intend to answer. Later in this chapter, we'll describe all five available models in the Sawtooth Software market simulator and provide some notes and recommendations for each.

After you have chosen the appropriate simulation technique, you can begin conducting simulations. Typically, one first examines the shares of preference (or choice) given to the products in the base case. Then, modifications to the base case are investigated by altering the base case itself and rerunning the analysis, or by adding additional "scenarios." A scenario is just another name for a defined set of competitive products, and setting up each subsequent scenario feels just like defining the first base case scenario. The market simulator lets you input many simulation scenarios, and stores these for your convenience.

Prior to introducing the different models of choice offered by the Market Simulator, it is instructive to cover two topics: the Exponent and the Red-Bus/Blue-Bus problem.

3.2   The Exponent (Scale Factor)

Assume that you set up a simulation as defined in the previous section with three products: A, B and C. Also assume that you are simulating the projected choice for just one individual under a Share of Preference model (described in greater detail later). After clicking **Compute!**, simulation results for this individual might come back as follows:

Product	Share of Choice
A	10.8%
B	24.0%
C	65.2%
Total	100.0%

Note that in conjoint simulations, the resulting shares are normalized to sum to 100%. We interpret these results to mean that if this respondent was faced with the choice of A, B, or C, he would have a 10.8% probability of choosing A, a 24.0% probability of choosing B, and a 65.2% probability of choosing C. Note that B is more than twice as likely to be selected as A, and C is more than twice as likely to be chosen as B.

Let's suppose, however, that the differences in share seen in this simulation are really greater than what we would observe in the real world. Suppose that random forces come to bear in the actual market (e.g. out-of-stock conditions, buyer confusion or apathy) and the shares (probabilities of choice) are really flatter. We can often tune the results of market simulations using an adjustment factor called the Exponent.

The table below shows results for the previous simulation under different settings for the Exponent:

Share of Choice under
Different Exponent Values

	0.01	0.5	1.0	2.0	5.0
Product A	33.0%	20.2%	10.8%	2.4%	0.0%
Product B	33.3%	30.1%	24.0%	11.6%	0.7%
Product C	33.7%	49.7%	65.2%	86.0%	99.3%
Total	100%	100%	100%	100%	100%

The exponent is applied as a multiplicative factor to all of the utility part-worths prior to computing shares (see Appendix A for details). As the exponent approaches 0, the differences in share are minimized, and preference is divided equally among the various product offerings. As the exponent becomes large, the differences in share are maximized, with nearly all the share allocated to the single best product. (Given a large enough multiplier, the approach is identical to the First Choice model, with all of the share given to a single product.)

If you have solid external information (such as existing market share data) and have reason to expect that conjoint shares should resemble market shares (see earlier assumptions), you may want to tune the exponent within simulations. Or perhaps you have choice data from a holdout choice scenario included in your conjoint survey. You may decide to tune the exponent so that simulated shares resemble these holdout shares. If you do not have solid external information, you probably should not change the exponent from the default value of 1.

Once the exponent is "set" for a data set, one typically does not change it from one simulation to the next.

3.3   The Red-Bus/Blue-Bus Problem

While market simulators have proven eminently useful for simulating buyer behavior, one of the most common simulation models (the Logit or Share of Preference model) has displayed a problematic result as characterized by the oft-cited Red-Bus/Blue-Bus problem. The underlying property leading to this problem is termed IIA, which is shorthand for "Independence from Irrelevant Alternatives." The basic idea of IIA is that the ratio of any two products' shares should be independent of all other products. This sounds like a good thing, and at first, IIA was regarded as a beneficial property.

However, another way to say the same thing is that an improved product gains share from all other products in proportion to their shares; and when a product loses share, it loses to others in proportion to their shares. Stated that way, it is easy to see that IIA implies an unrealistically simple model. In the real world, products compete unequally with one another, and when an existing product is improved, it usually gains most from a subset of products with which it competes most directly.

Imagine a transportation market with two products, cars and red busses, each having a market share of 50%. Suppose we add a second bus, colored blue. An IIA simulator would predict that the blue bus would take share equally from the car and red bus, so that the total bus share would become 67%. But it's clearly more reasonable to expect that the blue bus would take share mostly from the red bus, and that total bus share would remain close to 50%.

It is important to note that some degree of IIA is appropriate and useful within market simulations. In many markets, there is some degree of randomness to buyer behavior. It is not that people are irrational, but that buyers must balance the costs of making a utility maximizing decision against the costs of taking the time to make perfect decisions.   It is quite reasonable for rational buyers to make what on the surface may seem as haphazard decisions — especially for low-involvement purchases. A similar or even duplicate offering could thus be expected to capture more share in the real world than a rational simulation model might suggest.

In general, market simulation models based on disaggregate models of preference (utilities estimated at the individual level) are more immune to IIA difficulties than aggregate models of preference (aggregate logit, as offered by our CBC System). In addition to modeling respondent preferences at the individual level, there are market simulation methods that help deal with IIA. These are described in the next sections (details behind the calculations are provided in Appendix A).

3.4   Market Simulator Models

The Market Simulator offers five models:

1. First Choice
2. Share of Preference
3. Share of Preference with Correction for Similarity
4. Purchase Likelihood
5. Randomized First Choice

First Choice

This option is the simplest and is sometimes referred to as the "Maximum Utility Rule." It assumes the respondent chooses the product with the highest overall utility. The results for this option are invariant over many kinds of rescalings of the utilities. In particular, one could add any constant to all the levels for an attribute and/or multiply all part-worth utilities by any positive constant without affecting the shares for the simulated products.

The First Choice model requires individual-level utilities, such as those generated by ACA, CVA, ICE or CBC/HB. The First Choice model cannot be used with Latent Class or Logit runs for CBC.

The First Choice model is very intuitive and simple to implement. Its principal strength is its immunity to IIA difficulties (red-bus/blue-bus problem). In other words, the First Choice rule does not artificially inflate share for similar (or identical products). This property is especially important for product line simulations or situations in which some product offerings are quite similar to others in the competitive set.

Its principal weakness is that the share of preference results are generally more extreme than the other simulation models and one cannot adjust the steepness of the model using the exponent multiplier. We have seen evidence that the First Choice model's predictions can often be more extreme (especially when using CVA or ACA utilities) than market shares in the real world — especially for low involvement purchases.

Another weakness is that it reflects information only about the respondent's first choice. Information about the relative preference for the remaining products in the simulation is lost. As a result, standard errors for the First Choice model are generally higher than with the other models offered in the CCS. Sample sizes need to be larger for First Choice modeling than the other approaches to achieve equal precision of estimates.

We recommend using the First Choice model with ACA or CVA utilities if you have large sample sizes and have determined through holdout choice validation or, preferably, through validation versus actual market choices that the First Choice model accurately predicts shares better than the other approaches.

Share of Preference Models

The Share of Preference models (both with and without correction for product similarity) use the logit rule for estimating shares (see Appendix A for details). The product utilities are exponentiated and shares are normalized to sum to 100%.

The Share of Preference models result in "flatter" scaling of share predictions than the First Choice model. In general, we expect that this flatter scaling more closely matches what occurs in the real world. The Share of Preference models capture more information about each respondent's preferences for products than the First Choice method. Not only do we learn what product is

preferred, but we learn the relative desirability of the remaining products. This means that standard errors of share predictions are lower than the First Choice shares.

The Share of Preference model (without correction for product similarity) is subject to IIA, and can perform poorly when very similar products are placed in competitive scenarios (e.g. line extension simulations) relative to other less similar items within the same set. If using CBC under aggregate logit simulations, the IIA problem is intensified. Under Latent class, the problem is somewhat reduced. With individual-level utility models (ACA, CVA, ICE or CBC/HB), the problem is greatly reduced, but nonetheless can still be an issue.

The Share of Preference with Correction for Product Similarity model can result in more valid predictions when the competitive set includes products that have significant differences in similarities. However, this model is not as theoretically complete as the Randomized First Choice method and can give unexpected results, especially when conducting sensitivity simulations. The Randomized First Choice method has been shown to handle product similarity issues in conjoint simulations better. For this reason, we generally don't suggest using the Share of Preference with Correction for Product Similarity. It remains an option in the Sawtooth Software simulator mainly for historical purposes.

Purchase Likelihood Model

The purchase likelihood model estimates the stated purchase likelihood for products you specify in the simulator, where each product is considered independently. The likelihood of purchase projection is given on a 0 to 100 scale.

If you intend to use the Likelihood of Purchase option in the Market Simulator, your data must be appropriately scaled. The following estimation methods result in data appropriate for the purchase likelihood option:

1. ACA, if calibration concepts have been asked and used in utility estimation.

2. CVA, if single-concept presentation was used, and the logit rescaling option used with OLS regression.

3. ICE or CBC/HB, if calibration concepts have been asked and the CALIB program used to rescale the utilities.

Any other procedure will result in simulations that are not an accurate prediction of stated purchase likelihood. Also keep in mind that the results from the Purchase Likelihood model are only as accurate as respondents' ability to predict their own purchase likelihoods for conjoint profiles. Experience has shown that respondents on average exaggerate their own purchase likelihood.

You may use the Purchase Likelihood model even if you didn't scale the data using calibration concepts, but the results must only be interpreted as a relative desirability index. Meaning: a value of "80" is higher (more desirable) than a value of "60," but it doesn't mean that respondents on average would have provided an 80% self-reported likelihood of purchase for that particular product.

The purchase likelihoods that the model produces are not to be interpreted literally: They are meant to serve as a gauge or "barometer" for purchase intent. Under the appropriate conditions and discount adjustments (calibration), stated intentions often translate into reasonable estimates of market acceptance for new products.

Randomized First Choice

The Randomized First Choice (RFC) method combines many of the desirable elements of the First Choice and Share of Preference models. As the name implies, the method is based on the First Choice rule, and can be made to be immune to IIA difficulties. As with the Share of Preference model, the overall scaling (flatness or steepness) of the shares of preference can be tuned with the Exponent.

Most of the theory and mathematics behind the RFC model are nothing new. However, to the best of our knowledge, those principles have never been synthesized into a generalized conjoint/choice market simulation model. RFC, suggested by Orme (1998) and later refined by Huber, Orme and Miller (1999), was shown to outperform all other Sawtooth Software simulation models in predicting holdout choice shares for a data set they examined. The holdout choice sets for that study were designed specifically to include product concepts that differed greatly in terms of similarity within each set.

Rather than use the part-worth utilities as point estimates of preference, RFC recognizes that there is some degree of error around these points. The RFC model adds unique random error (variation) to the part-worth utilities and computes shares of choice in the same manner as the First Choice method. Each respondent is sampled many times to stabilize the share estimates. The RFC model results in a correction for product similarity due to correlated sums of errors among products defined on many of the same attributes.

The RFC model is very computationally intensive, but with today's fast computers speed is not much of an issue. With the suggested minimum of 100,000 total sampling iterations for a conjoint data set, it takes only a few moments longer than the faster methods to perform a single simulation. According to the evidence gathered so far on this model, we think it is worth the wait. The RFC model is appropriate for all types of conjoint simulations, based on either aggregate- or individual-level utilities.

The most complete use of the RFC model requires tuning the appropriate amount of attribute- and product-level error. By default, only attribute-level error is used in the simulator. This setting assumes no product share inflation for identical offerings. If you have questions regarding tuning the RFC model, please see Appendix A, confer with your consultant, or read the technical paper entitled "Dealing with Product Similarity in Choice Simulations," available for downloading from our home page: http://www.sawtoothsoftware.com.

3.5   External Effects

The factors evaluated in a conjoint analysis study usually focus on product attributes, ignoring other factors that affect market share. Consequently, the Market Simulator calculates a "preference share" rather than a "market share," emphasizing that important factors for calculating market share are missing from the model. Such missing factors include: the level and effectiveness of advertising, the size and effectiveness of the sales force, the number of outlets where the product is sold, the length of time the product has been on the market, and whether the product was first on the market.

The principal value of conjoint simulators is to indicate the kinds of changes that would make the most favorable differences in customer preferences. There is no real need for them to have the appearance of market shares. All the same, it is sometimes disconcerting for those viewing results to see shares of choice that are vastly different from known market shares.

The External Effect option of the Market Simulator helps to account for factors outside the model. When used properly, External Effects can lead to more realistic predictions. Even so, we recognize that the method used in our simulator to adjust for external effects is a simple approach with certain weaknesses. External Effects are introduced into the Market Simulator by applying a multiplicative External Effects factor to each product's preference share. The factor ranges from 0 to 9999. A value of 1 introduces no effect, a value greater than 1 increases a product's preference share above what the model would otherwise predict, a value less than 1 decreases the predicted preference below this value, and a value of 0 eliminates the product from the model.

You set External Effect factors from the **Scenario Specification** dialog, by checking *Apply External Effects.* When the *Apply External Effects* box is checked, an additional column in the product specification grid appears (at the far right) in which you can type external effects.

Setting the factor for each product is subjective. The following procedure minimizes the level of subjectivity:

1. Start by running a simulation on products currently on the market for which market shares are known or can be estimated. Set the External Effects factor to 1 for all products in this simulation.

2. Divide the actual market share for each product by the preference share predicted by the model. This number becomes the External Effects factor for that product.

3. Re-run the Market Simulator with these factors applied. Check that the results now reproduce the actual market shares for the products.

4. Set the External Effects factor for new products, using the External Effects factor for the current products as benchmarks. A new product should have an External Effects factor close to the ones for similar existing products (similar with respect to the "missing" factors listed above).

5. The Market Simulator produces shares that new products are expected to achieve when they have fully penetrated the market. The External Effects factor can be used to estimate shares before full penetration is achieved. One way to do this is to estimate, from past experience, the extent to which new products have reached full penetration at a time t after their introduction. When you want a simulation for time t, multiply this penetration factor by the one you have calculated using the steps above, and use this result as the External Effects factor.

For example, suppose you are a car manufacturer and you are introducing a new sedan. And suppose from your past sales you know the following penetration rate (fraction of expected monthly sales) for a sedan:

	Months After Introduction						
Time	0	3	6	9	12	15	18
Market Penetration	0.05	0.20	0.50	0.80	0.90	0.95	1.00

Then, if the External Effects factor at full market penetration is expected to be 1.2, the External Effects factor used as a function of time t would be:

	Months After Introduction						
Time	0	3	6	9	12	15	18
Market Penetration	0.06	0.24	0.60	0.96	1.08	1.14	1.20

As a final note, if you do need to apply external effect factors, we suggest you first investigate tuning the Exponent to best fit target shares prior to invoking external effect adjustments.

3.6   Interpolating between Levels

Sometimes you may want to simulate a product that has a level in between two levels you measured. This is called *interpolation*, and it is usually a reasonable thing to do when dealing with attributes that are quantitative in nature (price, speed, etc.). The market simulator uses a straight-line interpolation between adjacent part-worths.

Consider the following attribute:

Level #	Level Text
1	5 pages per minute
2	10 pages per minute
3	15 pages per minute
4	20 pages per minute

The simulator lets you specify a level code between two levels you measured. For instance, level 2.5 represents the point half-way between 10 and 15 pages per minute, or 12.5 pages per minute. Twelve pages per minute represents a value 40% of the way between 10 and 15 pages per minute. Therefore, level 2.4 corresponds to 12 pages per minute. But there is a much easier way to interpolate between levels. You may find it more convenient to refer to levels of quantitative attributes in the Market Simulator using values other than the original level numbers. When you click **Assign Level Values**, you can assign values to represent levels, such as:

Level #	Level Value	Level Text
1	5	5 pages per minute
2	10	10 pages per minute
3	15	15 pages per minute
4	20	20 pages per minute

It is much easier to specify products in the market simulator after recoding level numbers to values that match the pages per minute. To simulate a product at 12 pages per minute, we would now specify "12" for speed.

Three additional points are worth mentioning regarding recoding attribute levels for use in the simulator:

1. Interpolation only works correctly if all values are either ascending or descending for an attribute.

2. If you want to omit an attribute from simulations, you must omit it for all products in the simulation. You omit a level by specifying "N/A" (you type these characters) instead of a numeric code.

3. If linear terms (coefficients) were estimated instead of part-worths for a quantitative function such as price, you should make sure not to change the level values from those that your consultant provided. Please confer with your consultant for more information.

3.7   The "None" Weight

If you are conducting simulations for a CBC study and the questionnaire contains choice tasks including a "None" option, then the option of None can be included in simulations. However, the share of respondents predicted to choose None will only be correct if the number of products in the simulation is the same as the number of products in the CBC questionnaire. This is another problematic outcome of the IIA rule and especially affects simulators built upon aggregate logit runs. With individual-level modeling under CBC/HB, we have seen that the share of "None" is approximately correct irrespective of the number of products used in the simulator.

The market simulator lets you specify a "None" weight from ***Scenario Specification*** dialog by clicking the ***Advanced Settings...*** button. By default, the None weight is set to "0," which means that we do not report a None percentage and we assume that all respondents are in the market and must "buy" a product.

If you are using an aggregate logit run in your market simulations and are using a different number of products than were reflected in your CBC questionnaire, you may need to consider the suggestions below to deal with the problems of IIA and the "None" weight. If you are using individual-level utilities from CBC/HB, then you these considerations are less an issue for you.

None Calibration for Aggregate Logit Models

Logit (Share of Preference) models tend to give too much share to products that are similar to one another, and to penalize products that are unique (this is especially the case with aggregate logit solutions). The None option does not have specified levels on any of the conjoint attributes, and is therefore unique as compared to the products in the simulation.

If you do a simulation with a few products plus None, and then try another that includes those same products plus others, you will find that the share predicted for None will be smaller when there are more products. That may not be an accurate reflection of reality. If the respondents who chose None did so because they would *never* buy a product in that category, then it would clearly be incorrect to assume that fewer would choose none just because they are offered an array of more products. On the other hand, if respondents are really candidates for those products, then one would expect the share choosing None to decrease when respondents are offered a richer set of choices.

We know of no way to tell how the share choosing None should vary as the number of products in the simulation changes, although we think some allowance should be made in those cases where more products are in the simulation than in the original choice tasks.

For that reason, we provide the user with the capability of adjusting the percentage choosing None, through multiplication by a constant supplied by the user. In general, we think the None weight should be somewhere between unity and the fraction **n/t**, where **n** is the number of products in the simulation and **t** is the number of products in the average choice task. With individual-level utilities, less (or no) adjustment for the None may be needed. With aggregate utilities, more adjustment may be appropriate.

4 Chapter 4: Using the Market Simulator

4.1 The Market Simulator Dialog

Introduction

The previous two chapters have dealt mainly with theory. This chapter discusses how to use the software to simulate market choices. We should note that context-sensitive on-line help for all screens and dialogs within the software is available by pressing F1. You can access general help within the CCS (SMRT) by clicking **Help | Help Topics**.


The Market Simulator Dialog

To start the Client Conjoint Simulator, click **Start | Programs | Sawtooth Software | Sawtooth Software SMRT**. The individual who conducted the data analysis for your study should have given you five study files (see the list in Chapter 1). You should have copied these files into a directory on your hard drive (not a network directory).

The first step in running the Client Conjoint Simulator is to open a study. To open a study, select **File | Open**, and browse to the directory containing the studyname.SMT file. Double-click the studyname.SMT file. The main program menu is displayed.

To open the market simulator, click *Analysis | Market Simulator*. The Market Simulator dialog is displayed:
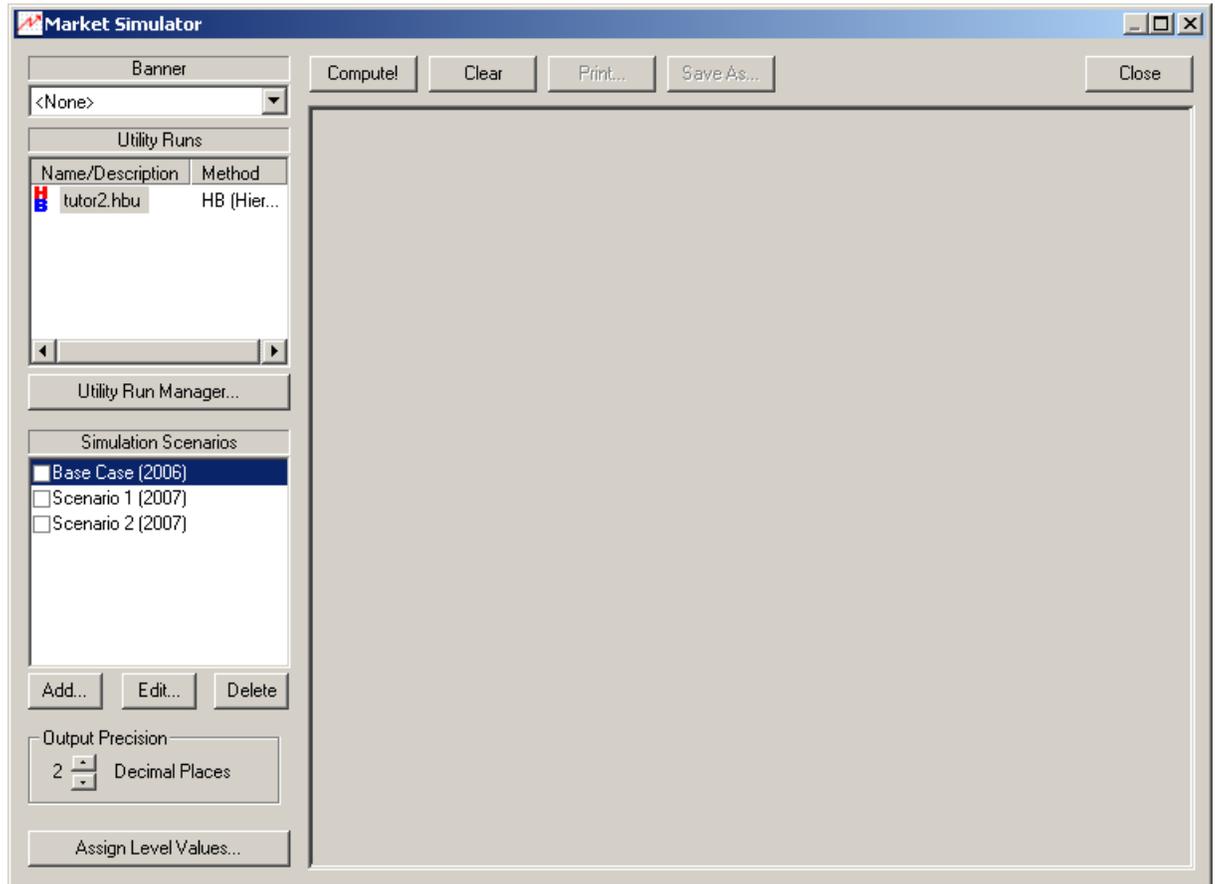
This dialog contains two main lists: the list of available Utility Runs (which you cannot modify) and the list of Simulation Scenarios (which you can add to and modify). A description of each of the settings and buttons is provided below:

Banner

If you are using part-worth utilities from a method other than logit, you can specify the banner variable to use in reporting the results of the simulation. For example, if you want to see the results split out between large and small business, you can choose that variable as the banner point.

Utility Runs

A utility run contains a set of part-worths utilities estimated in a particular way for a given number of respondents.   For example, perhaps your study uses CBC and your consultant used two multivariate methods to estimate respondent preferences: Logit and Latent Class.

One utility run may contain the solution from Logit and a second might contain the solution from Latent Class. Or, one utility run might just include the large business customers and another run the small business customers.

By default, the first utility run on the list is highlighted. Highlight the utility run you want to use during simulations by clicking with the mouse. Only one utility run can be selected at a time.

Simulation Scenarios   The consultant who delivered the data files to you may have initially specified some scenarios; alternatively, the list of simulation scenarios may be blank. To specify a new scenario, click **Add…**. To choose an existing scenario, check the scenario box with the mouse. You can check more than one simulation scenario at a time. If you choose multiple scenarios, all of the simulations will be conducted at the same time (in batch mode).

Output Precision   This setting controls how many decimal places of precision are used when reporting results.

Assign Level Values   This button lets you define the numeric code to be applied to different levels, such as for levels of price, speed or quantity. Assigning level codes can make it easier to define products in simulation scenarios and interpolate between measured levels. (See the section on interpolation in Chapter 3.)

Compute!   Click this to perform a simulation for the selected Utility run and Simulation Scenarios. The simulation results are displayed in the report window. If no products have been specified in a Simulation Scenario, average part-worth utilities and importances are displayed.

Clear   Click this to clear the results from the report window.

Print…   Prints the results displayed in the report window.

Save As…   Saves the results displayed in the report window to a text file.

Close   Exits the Market Simulator dialog.

4.2 Adding a Simulation Scenario

Usually the first step in using the market simulator is to define a "Base Case" scenario. A base case typically reflects a current (or future) market scenario: your brand vs. the relevant competition. If there is no relevant competition, or your conjoint study was designed to model only your product, the base case may be a single product.

The consultant who delivered the study files to use in the market simulator may already have defined a base case scenario or other scenarios. If so, when you open the Market Simulator dialog, these will appear in the list of Simulation Scenarios. You can view or edit an existing simulation scenario by highlighting that scenario and clicking *Edit…*.

To add a new scenario, click *Add…*. If a scenario or scenarios have already been defined, you are next asked whether the new scenario should be a copy of an existing scenario (used as a starting point, so that you don't have to re-enter common information again) or whether you want to start with a blank template.

The *Scenario Specification* dialog is then displayed:

The cursor is initially active in the *Name:* field. You must provide a name to identify the scenario. This name will appear in the Simulation Scenarios list in the *Market Simulator* dialog.

The Market Simulator lets you input the market scenario in a grid format, where the products are rows (you can have up to 100), and the attributes are columns (you can have up to 30). For example, you may have in mind to specify the following competitive scenario for a study that has

three total attributes with three levels each:

Product	Attribute 1 (Brand)	Attribute 2 (Performance)	Attribute 3 (Price)
Product A	1	2	3
Product B	2	1	2
Product C	3	3	1

If you started with a blank template, only one line in the grid is displayed. A default label of "Product 1" is provided, which you can edit by clicking this field with the mouse. To specify the attribute levels that make up a product, click on each cell within the grid and specify a level code. The attribute level labels and level codes are displayed automatically in the *Attribute Level View* window as an aid. You must provide a level value or an "N/A" for each cell in the grid under the attribute columns.

After you have fully defined the first product (row) of the grid, you are ready to add another product. You can add and delete products from the simulation scenario using the ***Insert Product*** and ***Delete Product*** buttons.

Insert Product   The products in the market scenario are specified in the grid central to this dialog. Each row of the grid represents a product concept. When you click ***Insert Product***, a new row is added to the grid. You define a product concept by typing in level values directly into the grid for the attributes describing that product.

Delete Product   Deletes the highlighted product concept.

OK   Saves any modifications you've made to the current dialog and returns you to the main market simulation dialog.

Cancel   Returns you to the main market simulation dialog without saving any changes.

4.3 Other Controls in Simulation Scenarios

Besides formulating the different competing products in the simulation scenario, there are a number of other important decisions to be made. You must decide on the simulation method (the mathematical model that converts respondent preferences into simulated shares of choice or purchase likelihood estimates), respondent settings (if using individual-level data), the scaling of shares and output display options.

Simulation Method   Select the method used for simulations for the current scenario. The method is Randomized First Choice by default. (See Chapter 3 for more details about the different simulation methods.)

Operating Mode   The default mode is *Simulation*. You can specify *Sensitivity* mode, which is a way of processing in batch multiple simulation runs that are variations of the current scenario. For example, you may wish to see how varying all levels of price for Product 1 affects its share relative to a fixed set of competitors. Rather than run the same simulation over and over again for all levels of price for Product 1, you can specify *Sensitivity* mode, and choose Product 1 as the *Sensitivity Product*, and Price as the *Sensitivity Attribute*.

Respondents to Include   If you are using part-worth utilities from a method other than logit, you can specify the respondents to include in the simulation. If you are using aggregate utilities (from Logit), this option is not available.

Respondent Weights   If you are using individual-level utilities (from ACA, CVA, ICE or CBC/HB), you can specify the respondent weights to be applied in the simulation. If you are using aggregate utilities (from Logit or Lclass), this option is not available.

Apply External Effects   In some situations, you may want to adjust the share given to a particular product by some external factor. For example, perhaps that product is only available in half of the markets relative to the other products. In that case, you can check the *Apply External Effects* box. When this box is chosen, a new column appears in the product specification grid labeled "External Effect." All external effects are initially set to unity (1). To adjust the shares to account for a given product only being available in half of the markets, you might specify an external effect of 0.5 for that product, and leave all other products with external effects of unity (1).

Advanced Settings   When you click this button, a dialog appears that lets you define additional settings:

Correlation Cutoff   If you are using individual-level utilities generated by ACA, CVA, ICE or CBC/HB, a measure of fit is provided indicating the relative consistency of each respondent's utilities. You can select a value used as a cutoff for excluding respondents that are at or below that threshold for consistency.

Exponent   The exponent tunes the overall "flatness" or "steepness" of the share results (see discussion in Chapter 3). It has effect only for the Share of Preference models and Randomized First Choice.

"None" Weight   If you are analyzing data for a CBC study that included a "None" option, this lets you specify the weight applied to the None utility. By default, it is zero (no "None" share computed). (See

discussion in Chapter 3 for more direction.)

Unacceptable   If you are using data from an ACA study and if you asked the
Extrapolation   "Unacceptables" section, you can tune how low the utility should
                          be for unacceptable levels. This is an an advanced setting and
                          we suggest you confer with your consultant regarding the proper
                          setting to use in this case.

Output Options

Individual Results to File

If you are running simulations using part-worth utilities generated by a method
other than logit, you can save the estimated shares of preference for each
individual to an ASCII file for analysis in another software package. This can
be helpful, for instance, if you want to identify particular respondents that
prefer (or do not prefer) a certain product.

Display Utilities

Controls whether the utilities for each attribute level (or interaction terms) are
displayed in the output. The utilities displayed are rescaled by a method called
Zero-Centered Diffs. The diffs method rescales utilities so that the total sum
of the utility differences between the worst and best levels of each attribute
across attributes (main effects) is equal to the number of attributes times 100.
Note: the average part-worth attribute utilities are influenced by the number of
respondents in the simulation and respondent weighting, but are not affected
by the product specifications you enter. For more information about
interpreting conjoint utilities, please refer to Chapter 6.

Display Importances

Checking this option tells the CCS to include a summary of attribute
importances in the simulation output. The importance of an attribute is
defined as its weight, or the maximum influence it can have on product choice,
given the range of attribute levels defined in the study. In other words, an
attribute's importance is an indicator of the amount of influence an attribute
may have in simulations. Attribute importances are only accurate when
individual- or segment-level utility data are available (such as computed under
Latent Class, ACA, CVA, ICE or Hierarchical Bayes). Aggregate importances
from logit can be misleading if respondents disagree about the order of
preference of levels within an attribute. Please see Chapter 6 for more
information on attribute importance computations and interpretation.

4.4   Running the Simulation

After you have specified the products and chosen the desired settings for your simulation scenario, you are ready to compute the results. From the *Market Simulator* dialog, click *Compute!* Depending on the model you choose and the number of respondents, it can take from a few moments to a few minutes to compute. The results are displayed in the text report window. You can scroll up and down across the report within this window. You can also cut-and-paste results into a word processor or spreadsheet package.

Every time you click *Compute!*, results are appended to any existing data in the report window. Click *Clear* to delete the contents of the window.

5   Chapter 5: Practical Simulation Examples

5.1   Introduction

In this section, we'll discuss how to use the market simulator to answer some common strategic questions introduced in Chapter 2:

- Given a current competitive environment, what product should I offer to maximize interest in my offering?

- What is the relative price sensitivity of different brands? If I raise my price by 10%, how will it affect my brand?

- What portfolio of products can I offer to appeal to different market segments and maximize overall share?

For the examples in this chapter, we'll assume the following three attributes, measured using three levels each:

Brand:
1   BrandA
2   BrandB
3   BrandC

Style:
1   StyleA
2   StyleB
3   StyleC

Price:
1   $100
2   $150
3   $200

The results reported in this chapter are fictitious and not based on any particular data set.

5.2   New Product Introductions

Let's assume that your company is interested in entering a market that currently consists of just two competitors. There are just three attributes that adequately describe the products and account for preference in the market: Brand, Style and Price. The two players are:

    1) BrandA, StyleA, $100
    2) BrandB, StyleB, $200

Your company has developed a new style (StyleC) that you think may appeal to buyers, and you want to investigate its potential with respect to the two existing products.

The first step, typically, is to simulate the existing market scenario. You use the market simulator to define the two existing products:

Product Specifications:

Product Name	Brand	Style	Price
"BrandA"	1	1	1
"BrandB"	2	2	3

For this simulation scenario, there are no similarities (in terms of common shared attribute levels) between the products in the simulation. Therefore, the Share of Preference model is appropriate (though it would also be appropriate to use the Randomized First Choice method). When you click **Compute!**, the following shares of preference are displayed:

Shares of Preference for Products:

BrandA	64.3
BrandB	35.7

Note that the buyers in the simulation are all assumed to choose a product, so the shares of preference across products in the simulation sum to 100%.

Let's assume that you have actual market share information about these two brands. You note that the shares reported above do not necessarily match the actual market shares. You accept this, however, recognizing that many factors influence market shares in the real world that cannot be captured through conjoint analysis. You are principally interested in *relative* preferences, assuming that the marketplace is an equal playing field: equal distribution, awareness, effectiveness of sales force, and equilibrium long-range demand.

In the second stage of this simulation example, we'll define a new scenario that includes your company's proposed product: BrandC, StyleC, $150. You add another row to the simulation specification grid:

Product Specifications:

Product Name	Brand	Style	Price
"BrandA"	1	1	1
"BrandB"	2	2	3
"BrandC"	3	3	2

After clicking **Compute!**, the following shares are displayed:

Shares of Preference for Products:

BrandA	42.5
BrandB	21.3
BrandC	36.2

You note that BrandA is still the most preferred product, but that your brand is preferred to BrandB.

Like any market research statistic computed from samples, shares of preference are not estimated without error. It is common to estimate a confidence interval, to gain a feel for the degree of uncertainty due to sampling and measurement error associated with a given share of preference. If your simulator is based upon part-worth utilities generated by a method other than logit, the CCS lets you add the standard errors of the shares of preference to the simulation report. Let's assume that the standard error reported for BrandC for the simulation above was 1.53. The 95% confidence interval is computed by adding plus and minus 1.96 times the standard error to the estimated share of preference. In this example, the 95% confidence interval is 36.2 plus and minus (1.96)(1.53) = 3.0 share points, or the interval [33.2, 39.2].

You next may ask yourself what price you would need to charge to capture the same relative preference as BrandA. To simulate this, you lower the price slightly for your brand. The market simulator lets you interpolate between levels, so you can investigate even the smallest of price changes. As a first step, you decide to lower the price to $130 for BrandC (while holding the specifications for BrandB and BrandA constant). The new simulated shares are:

Shares of Preference for Products

BrandA	39.2
BrandB	19.0
BrandC	41.8

You have overshot the mark (BrandC's share exceeds BrandA's share), so you try a slightly higher price than $130 and click **Compute!** once more. You make repeated attempts until BrandA and BrandC's shares are equal. Let's assume that after a few more attempts, you discover that the price that makes your company's offering match the share of preference of the market leader is $136. Another way of thinking about this finding is that your proposed product commands a $136 - $100 = $36 premium over BrandA's product. (Respondents are indifferent between BrandA, StyleA at $100 and BrandC, StyleC at $136).

5.3   **Estimating Demand Curves and Elasticities**

Before outlining simulation strategies for pricing research, we should give a few notes of caution. The ACA system is often not a very good pricing research tool; CVA and particularly the CBC system are considered stronger approaches for most markets.

If using CBC for pricing research, we recommend the use of Latent Class or especially individual-level computation methods, such as the CBC/HB (Hierarchical Bayes) module.

For estimating demand curves using market simulations, the Share of Preference method is appropriate. Randomized First Choice may also be selected, if you believe some correction for product similarity should result if products carry the same prices (however, the use of Randomized First Choice for generating demand curves can often result in strange shifts in the curves around the "average" price, if brands are initially set at the same average price).

We will build upon the previous example during this section. We previously computed shares of preference for three products defined using the following attribute level codes:

Product Specifications:

Product Name	Brand	Style	Price
"BrandA"	1	1	1
"BrandB"	2	2	3
"BrandC"	3	3	2

You may recall that price levels 1, 2 and 3 correspond with the following prices $100, $150, and $200. The shares of preference for the products as defined above were:

Shares of Preference for Products

BrandA	42.5
BrandB	21.3
BrandC	36.2

Let's assume we wanted to estimate a demand curve for your company's offering: BrandC, in the context of the current competition and prices. We do this through "sensitivity analysis." The CCS offers an automatic way to conduct sensitivity analysis, but so you understand the process, we'll describe it in separate steps.

Recall that we measured three distinct levels of price $100, $150 and $200. Note that we've already computed the share of preference for BrandC when it is offered at $150 (21.3).

To estimate the demand curve for BrandC, we'll need to conduct two additional simulations: a simulation with BrandC at the lowest price ($100), and a simulation with BrandC at the highest price ($200). For each of these simulations, we'll hold the BrandA and BrandB product specifications constant.

To estimate BrandC's share at the lowest price ($100), we use the following product specifications:

Product Specifications:

Product Name	Brand	Style	Price
"BrandA"	1	1	1
"BrandB"	2	2	3
"BrandC"	3	3	1

We click **Compute!**, and the following shares are reported:

Shares of Preference for Products:

BrandA   33.9
BrandB   15.6
BrandC   50.5

We record BrandC's share (50.5), and proceed to the next step. To estimate BrandC's share at the highest price ($200), we use the following product specifications:

Product Specifications:

Product Name	Brand	Style	Price
"BrandA"	1	1	1
"BrandB"	2	2	3
"BrandC"	3	3	3

We click **Compute!**, and the following shares are reported:

Shares of Preference for Products

BrandA   49.2
BrandB   26.9
BrandC   23.9

From these three separate simulation runs, we now have all the information we need to plot a demand curve for BrandC, relative to the existing competitors and prices. Assuming that BrandA and BrandB are held constant at current market prices, the relative shares of preference for BrandC at each of the price points within the measured price range are:

Price	Share of Preference
$100	50.5
$150	36.2
$200	23.9

We have demonstrated how to estimate a demand curve for BrandC, relative to the existing competitors at current market prices. If the goal is to estimate demand curves for all brands in the study, the usual procedure is to record the share for a brand at each price level while holding all other brands at the *average* (middle) price. It is often interesting to plot these demand curves and look at the patterns of price sensitivity between brands and the different slope of the curves from one segment of the curve to the next. It is also common to want to characterize the degree of price sensitivity using a single value, referred to as an elasticity. The *Price Elasticity of Demand* (E) is defined as:

$$E = \frac{\%\triangle Q}{\%\triangle P}$$

where,
    "%△" means "percent change in,"
    Q is defined as quantity demanded, and
    P refers to the price.

If the brand or product follows the law of supply and demand (most products do), price increases lead to decreases in quantity demanded, and the elasticity is negative. The larger the absolute value of the elasticity, the more price sensitive the market is with respect to that brand or product.

Using the "midpoints" formula, we can compute the average price elasticity of demand across the demand curve for BrandC:

$$E = \frac{q2 - q1}{(q1+q2)/2} \qquad + \qquad \frac{p2 - p1}{(p1+p2)/2}$$

$$\frac{23.9 - 50.5}{(50.5+23.9)/2} \qquad + \qquad \frac{200 - 100}{(100+200)/2}$$

$$-0.715 \qquad + \qquad 0.667$$

$$= -1.073$$

Another way to compute the average price elasticity of demand (which can be more accurate if more than two price points along the curve have been estimated) is the "log-log" regression. One takes the natural log of prices and shares and regresses the log of share on the log of price (you can do this within a spreadsheet). The resulting beta is the average price elasticity of demand.

As with all conjoint simulation results, the resulting elasticities from conjoint simulators are interpreted bearing in mind some assumptions.   In particular, the degree of noise within the conjoint data is particularly relevant. For example, if the respondents to the conjoint survey answered in a more haphazard way compared to buyers in the real world, the price elasticities estimated from conjoint simulations may be uniformly understated (too insensitive). Also, you should recognize that changing the Exponent will shift the elasticities. Even if this is the case, the relative price sensitivities for brands are still useful.

5.4   Designing Products to Appeal to Unique Market Segments

Customizing products to appeal to target segments or even individuals is a common theme in marketing. Many companies dedicate significant resources to developing a portfolio of products that it hopes will appeal to unique segments. For line extensions, the challenge for any company is to design new product(s) that take share from its competitors without stealing an unacceptable amount of share from products within its existing line.

One common approach to designing an effective line extension is to use the conjoint data to segment the market into latent (not observed) market segments (sometimes referred to as clusters) that have similar preferences. These segments are termed "latent," because they are not simply delineated based on an explicit variable such as gender, income or company size. Rather, the underlying segments are revealed through a statistical segmentation technique such as cluster analysis or latent class modeling. Segments are formed with the goal of maximizing the differences in preference between groups while minimizing the differences in preference within groups. Once these latent segments have been identified, one can profile them in terms of other variables in the survey (i.e. demographics, usage or media habits).

If the analyst who delivered the CCS to you used a Latent Class (for CBC data) or a cluster technique, he/she may have provided a way for you to conduct simulations "by segment." If using a Lclass utility run, this happens automatically. If based on a cluster analysis, the segment membership information can be made available to you as a "banner point."

For example, let's assume that a cluster analysis revealed three relatively different segments for the hypothetical example we've been using throughout this chapter. Let's also assume that the analyst who prepared the CCS for you made the segment membership information available to you as a "banner" variable called "SEGMENT."   If you select SEGMENT as the banner variable and click ***Compute!,*** the market simulator displays the part-worth utilities and importances for each segment.

By examining the part-worths and importances (see Chapter 6) for each group, you can gain insight into the product features that might appeal to each. You also should bear in mind the size of each segment, as this represents its demand potential. Consider the following part-worth utility preferences:

Part-Worth Utility Preferences

	Segment1 n = 128	Segment2 n = 283	Segment3 n = 216
BrandA	39	-51	-44
BrandB	5	39	-29
BrandC	-44	12	73
StyleA	61	-52	-34
StyleB	-23	45	-9
StyleC	-38	7	43
$100	56	55	50
$150	7	2	6
$200	-63	-57	-56

We can study the part-worths to learn about the differences among the segments. We can also use these preferences to simulate market choices for the following market scenario:

Product Specifications:

Product Name	Brand	Style	Price
"BrandA"	1	1	1
"BrandB"	2	2	3
"BrandC"	3	3	2

Shares of Preference for Products, by Market Segment:

	Segment1 n = 128	Segment2 n = 283	Segment3 n = 216	Total 627
BrandA	84.8	21.5	22.2	34.7
BrandB	7.4	40.0	14.2	24.5
BrandC	7.8	38.5	63.6	40.8

(Note that these shares do not match the shares reported for earlier examples in this chapter. Since these results are for illustration only, no significance should be attached to this difference.)

Let's assume your company produces BrandC with StyleC at $150. Your total share of preference is 40.8%. We see from the simulation by segment that yours is the most preferred product within Segment 3, and the second-most preferred product in Segment 2. BrandA clearly dominates Segment1 (the smallest segment).

Let's assume that your company was interested in offering an additional product. We could examine the table of part-worth utilities presented earlier as a first step in formulating hypotheses about what additional product might be successful.

Starting in order, you may first consider Segment 1, but this segment does not seem to offer many opportunities for your brand. BrandA offering StyleA at a low price has got this relatively small segment nearly wrapped up, and this segment doesn't seem very receptive to your brand: BrandC.

You next consider Segment 2, which seems to represent a better opportunity for your brand. It is a relatively large segment that prefers BrandB, but also seems receptive to BrandC. Note also that Segment 2 strongly prefers StyleB, but your company currently offers only StyleC. By offering a StyleB product, you might be able to convert some current BrandB customers from within Segment 2 to your product line.

You currently dominate Segment 3 and should probably not consider designing another product to appeal to this segment, since a good deal of the possible share to be gained from a new product would be taken from your existing product within that segment.

Let's simulate what happens if in addition to your current product (BrandC, StyleC, $150), you offer another product (BrandC, StyleB, $200). Note that since this simulation involves products that have similar definitions (there are two BrandC products), you should probably use a technique that adjusts for product similarities, such as Randomized First Choice.

Shares of Preference for Products, by Market Segment

	Segment1 n = 128	Segment2 n = 283	Segment3 n = 216	Total 627
BrandA	82.2	17.2	18.6	31.0
BrandB	7.2	32.0	11.9	20.0
BrandC (old)	6.8	27.7	47.8	30.4
BrandC (new)	3.8	23.1	21.7	18.7

The new product has somewhat cannibalized the existing product, reducing its share from 40.8 (see the previous simulation) to 30.4, but has resulted in a relative overall gain of 1 - [(30.4 +

18.7)/40.8] - 1= 20% in preference.

For line extension simulations you conduct, the answer will likely not be so clear and the process so direct as we've shown here. You'd certainly want to investigate other product configurations to make sure you weren't overlooking even better opportunities to enhance share. You would also want to consider the cost implications of different options for line extensions. Also, you would probably want to conduct sensitivity analysis for the new product with respect to price, to determine a strategic price point (given your costs and market share goals).

Viewing the preferences and shares by segment is not required in designing an effective line extension. However, viewing the separate market segments can help you more quickly recognize patterns of preference, size the different segments of the market, and thus more easily arrive at a good solution.

This exercise of viewing segment-based preferences and designing products to fill heterogeneous needs is a useful approach. However, it would seem more efficient to let an automated search algorithm find an optimal product or set of products rather than to proceed manually. Sawtooth Software's SMRT Advanced Simulation Module includes automated optimization search capability.

6   Chapter 6: Interpreting Conjoint Results

6.1   The Nature of Quantitative Data

Introduction

(The following is adapted from two separate articles originally published in our newsletter, *Sawtooth Solutions.*)

Our conjoint software systems provide a number of outputs for analyzing results including: utilities, importances, shares of preference and purchase likelihood simulations. This chapter discusses these measures and gives guidelines for interpreting results.

Before focusing on conjoint data, it may be helpful to review some fundamentals for interpreting quantitative data. The definitions below are adapted from *Statistics for Modern Business Decisions*, Fourth Edition, by Lawrence L. Lapin.

There are four general types of quantitative data:

1. **Nominal data** are those wherein the numbers represent categories, such as 1=Male, 2=Female; or 20=Italy, 21=Canada, 22=Mexico. It is not appropriate to perform mathematical operations such as addition or subtraction with nominal data, or to interpret the relative size of the numbers.

2. **Ordinal data** commonly occur in market research in the form of rankings. If a respondent ranks five brands from best "1" to worst "5," we know that a 1 is preferred to a 2. An example of an ordinal scale is the classification of the strength of tornados. A category 3 tornado is stronger and more damaging than a category 2 tornado. It is generally not appropriate to apply arithmetic operations to ordinal data. The difference in strength between a category 1 and 2 tornado is not necessarily equal to the difference in strength between a category 2 and a 3. Nor can we say that a category 2 is twice as strong as a category 1 tornado.

3. **Interval data** permit the simple operations of addition and subtraction. The rating scales so common to market research provide interval data. The Celsius scale also is interval scaled. Each degree of temperature represents an equal heat increment. It takes the same amount of heat to raise the temperature of a cup of water from 10 to 20 degrees as from 20 to 30 degrees. The zero point is arbitrarily tied to the freezing point of distilled water. Sixty degrees is not twice as hot as 30 degrees, and the ratio 60/30 has no meaning.

4. **Ratio data** permit all basic arithmetic operations, including division and multiplication. Examples of ratio data include weight, height, time increments, revenue and profit. The zero point is meaningful in ratio scales. The difference between 20 and 30 kilograms is the same as the difference between 30 and 40 kilograms, and 40 kilograms is twice as heavy as 20 kilograms.

6.2   Conjoint Part-Worth Utilities

Part-worth utilities characterize the relative desirability or attractiveness of attribute levels: the higher the utility, the more desirable it is to respondents. Conjoint utilities from ACA, CVA, CBC's Logit, Latent Class, ICE and CBC/HB are scaled to an arbitrary additive constant within each attribute and are **interval data**. The arbitrary origin on the scaling within each attribute results from dummy coding in the design matrix. For example, in CBC, logit utilities are scaled to sum to 0 within each attribute. A plausible set of utilities for miles per gallon might look like:

30 MPG	-100
40 MPG	0
50 MPG	100

Just because 30 MPG received a negative utility value does not mean that this level was unattractive. In fact, 30 MPG may have been very acceptable to all respondents. But, relatively speaking (all else being equal) 40 MPG and 50 MPG are better. The utilities are scaled to sum to 0 within each attribute, so 30 MPG must receive a negative utility value. (ACA, CVA and calibrated ICE utilities are also scaled to an arbitrary additive constant, but the sum of attribute utilities is not always equal to 0.)

Users of conjoint analysis are sometimes confused about how to interpret utilities. Difficulty most often arises in trying to compare the utility value for one level of an attribute with a utility value for one level of another attribute. With most conjoint methods, including Choice-Based Conjoint, it is not correct to compare a single value for one attribute with a single value from another. Instead, one must compare differences in values.

The following example featuring zero-centered conjoint utilities illustrates this point:

Brand A	40	Red	20	$ 50	40
Brand B	-60	Blue	0	$ 75	0
Brand C	20	Pink	-20	$100	-40

It is not correct to say that Brand C has the same desirability as the color Red. However, it is correct to conclude that the difference in value between brands A and C (40 - 20 = 20) is the same as the difference in values between Red and Blue (20 - 0 = 20). This market should be indifferent on average between Brand A in a Blue color (40 + 0 = 40) and Brand C in a Red color (20 + 20 = 40).

Even though we are comparing utilities within the same attribute, we cannot say that Brand A is twice as preferred as Brand C (40/20). Interval data do not support ratio operations.

6.3   Conjoint Importances

Sometimes we want to characterize the relative importance of each attribute. We do this by considering how much difference each attribute could make in the total utility of a product. That difference is the range in the attribute's utility values. We percentage those ranges, obtaining a set of attribute importance values that add to 100, as follows:

		Range	Percent Importance
Brand (A - B)	40 - -60 =	100	45.4%
Color (Red - Pink)	20 - -20 =	40	18.2%
Price ($50 - $100)	40 - -40 =	80	36.4%
		220	100.0%

The importance of Brand is 45.4%, the importance of Color is 18.2%, and the importance of Price is 36.4%. Importances depend on the particular attribute levels chosen for the study. For example, with a narrower range of prices, Price would have been less important.

We caution against calculating attribute importances from average conjoint utilities or aggregate logit analysis. When summarizing attribute importances for groups, it is best to compute importances for respondents individually (such as from CBC/HB analysis) and then average them, rather than computing importances using average utilities. Unless you compute importances at the individual level, you risk understating the importance of attributes about which there was disagreement regarding which level was preferred. For example, suppose we were studying two brands, A and B. If half of the respondents preferred each brand, the average utilities for A and B would be tied, and the importance of brand would appear to be zero!

Importances are **ratio data**. An attribute with an importance of 20 (20%) is twice as important as an attribute with an importance of 10.

Although importance scores have been commonly used over the last few decades, they are perhaps not as intuitive and meaningful as results from sensitivity analysis. Importance scores consider only the extreme utilities within an attribute, irrespective of whether the part-worth utilities actually follow rational order. Importance computations take advantage of even random variation in the utilities. Such problems are absent when characterizing the impact of attributes through sensitivity analysis.

6.4 Shares of Preference

All of our conjoint systems offer share of preference simulations. When two or more products are specified in the market simulator, we can estimate what percent of the respondents would prefer each product. Shares of preference are ratio data. Even so, we recognize that both the exponent (scaling multiplier) and the simulation model used can dramatically affect the scaling of shares of preference. A product which captures twice as much share as another in a first choice simulation (or using a large exponent) may capture considerably less than twice the share using the share of preference (probabilistic) model.

6.5   Purchase Likelihoods

During ACA, CVA or CBC interviews, respondents may be asked to rate individual products on a 0 to 100 point purchase likelihood scale. This is very helpful to gauge respondent interest in the product, and for scaling the data for use in purchase likelihood simulations. Once we have scaled conjoint data to reflect purchase likelihoods (see Appendix A), we can predict how respondents would have rated any combination of attributes included in the study in terms of that same measure.

Purchase likelihoods should not be considered as strictly **ratio data**. A respondent may not *truly* be twice as likely to purchase a product he rated a 50 versus another he rated a 25. Even so, it is quite common to state that a product with a purchase likelihood of 55 represents a 10% relative increase in purchase likelihood over a product which received a 50.

7   Appendix A: Technical Details

7.1   First Choice Method

This option is the simplest and is often referred to as the "Maximum Utility Rule." It assumes the respondent will choose that product with the highest overall utility. All other less preferred products receive zero share for this individual. The results for the First Choice model are invariant over many kinds of rescalings of the utilities.

If two or more products are exactly the same in terms of utility (and more preferred than any other product in the simulation scenario), the respondent's choice is split evenly amongst them.

7.2   **Share of Preference Options**

These options do not assume that the respondent always chooses the product with highest utility. Instead, they estimate probability of choosing the simulated product, arriving at a "share of preference" for the product.

This is done in two steps:

1.   Subject the respondent's total utilities for the product to the exponential transformation (also known as the antilog): s = exp(utility).

2.   Rescale the resulting numbers so they sum to 100.

Suppose two products, A and B, have total utilities 1.0 and 2.0. Then their shares of preference would be computed as follows:

	product utility	exp(utility)	share of preference
A	1.0	2.72	26.9
B	2.0	7.39	73.1

(For the Share of Preference with Correction for Product Similarity, shares of preference are modified by further computations.)

Unlike the First Choice option, the scaling of utilities can make a big difference with the Share of Preference Options. Consider what happens if we multiply the utilities by constants of 0.1 or 10.0:

For multiplier of 0.1

	product utility	exp(utility)	share of preference
A	1.0	1.105	47.5
B	2.0	1.221	52.5

For multiplier of 10.

	product utility	exp(utility)	share of preference
A	10.0	22026	00.005
B	20.0	4.85E08	99.995

If we multiply the utilities by a small enough constant, the shares of preference can be made nearly equal. If the utilities are made small enough, every product receives an identical share of preference, irrespective of the data.

If we multiply the utilities by a large enough constant, the shares of preference can be made equivalent to the First Choice model.

It is apparent that scaling of the utilities can make a big difference with the Share of Preference Models. Most of Sawtooth Software's utility estimation methods result in utilities appropriate for use with Share of Preference models. We strongly suggest you include holdout (fixed) choice tasks in your conjoint questionnaires to use for calibration, or check the simulation predictions against actual market share.

The Market Simulator lets you scale the utilities within the Share of Preference option at the time the simulation is done. This is accomplished by a parameter called the "exponent" that you can set when preparing for simulations. The default value of the exponent is 1. It has the same effect as the multiplier illustrated immediately above. The exponent can be used to adjust the sensitivity of the simulation results so that it more accurately reflects holdout choices, or actual market behavior.

A smaller exponent causes small shares to become larger, and large shares to become smaller — it has a "flattening" effect. In the limit, with a very small exponent every product receives the same share of preference.

A large exponent causes large shares to become larger, and small shares to become smaller — it has a "sharpening" effect. In the limit, a very large exponent produces results like those of the First Choice option.

If you have solid external information (such as existing market share data) and have reason to expect that conjoint shares should resemble market shares (see assumptions listed in Chapter 2), you may want to tune the exponent within simulations. If you do not have solid external information, you probably should not change the exponent from the default value of 1.

Share of Preference Model with Correction for Product Similarity

Share of preference models that do not consider similarities among products have a serious problem: if an identical product is entered into the simulation twice, it can receive up to twice as much total share of preference as it would when entered only once. Although no researcher would make the mistake of entering the same product twice, the principle is still troublesome. Products differ in similarity to one another, and "plain" share of preference models tend to give too little share to relatively unique products.

The third choice model in the Market Simulator includes a correction to prevent the preference shares of similar products from being overstated. The correction is based on each product's total similarity with other products. The basic idea behind this correction was suggested by Richard D. Smallwood of Applied Decision Analysis, although we are responsible for the details of its implementation.

The procedure is:

1. For the n products in a simulation, an n x n similarity matrix is computed, with 1's indicating complete similarity, 0's indicating total lack of similarity, and fractional values for differing degrees of similarity:

    First a "dissimilarity" matrix is computed. Consider a scale for each attribute where its levels are coded 10, 20, 30, and so on. The dissimilarity of a pair of products on an attribute is taken as the absolute difference between their codes, but with a maximum of 10. (This allows "continuous" and "categorical" attributes to be treated in the same way.) The total dissimilarity between two products is the sum of their dissimilarities over all attributes. Two products differing by an entire level on each attribute are maximally dissimilar.

    Next, total dissimilarities are rescaled by a constant so the maximum possible is 3.0, rather than 10 times the number of attributes.

    Dissimilarities are then converted to similarities by a negative exponential transformation. At this point the minimum possible similarity is exp(-3) ~ .05, achieved if two products have completely different levels on every attribute, and the maximum possible similarity is exp(0) = 1, achieved only if they have identical levels on all attributes.

    Finally, the similarities are subjected to a further rescaling that sets the minimum to 0 and the maximum to 1.

2. Column totals of the similarity matrix are calculated to get a vector of "total similarities." The smallest possible value is 1.0, which is found if a product is maximally dissimilar to all others and similar only to itself. If two products are identical but maximally dissimilar to all others,

those products each have values of 2.0.

3. For each respondent, shares of preference are divided by corresponding "total similarities," and then renormalized to have sum of unity. This has the effect of reducing shares for products that are relatively similar to others, and increasing shares for products that are relatively unique. In the limit, where two or more products are identical and totally unlike any others, they divide among themselves the share that each product would have if the others were not present.

7.3   Purchase Likelihood Option

The purchase likelihood option estimates the stated purchase likelihood for products you specify in the simulator. Each product is considered independently. If you intend to use the Likelihood of Purchase option in the Market Simulator, your data must be appropriately scaled. The following estimation methods result in data appropriate for the purchase likelihood option:

1.   ACA, if calibration concepts have been asked.

2.   CVA, if single-concept presentation was used, and the logit rescaling option used with OLS regression.

3.   ICE or CBC/HB, if calibration concepts have been asked and the CALIB program used to rescale the utilities.

Any other procedure will result in simulations that are not an accurate prediction of stated purchase likelihood. Also keep in mind that the results from the Purchase Likelihood model are only as accurate as respondents' ability to predict their own purchase likelihoods for conjoint profiles. Experience has shown that respondents on average exaggerate their own purchase likelihood.

You may use the Purchase Likelihood model even if you didn't scale the data using calibration concepts, but the results must only be interpreted as a relative desirability index.

The purchase likelihoods that the model produces are not to be interpreted literally: They are meant to serve as a gauge or "barometer" for purchase intent.

This model provides a means of simulating a product category with only a single product. The other three choice models are comparative models that require at least two products to be specified in a simulation. Likelihoods are estimated for product concepts by summing scaled utilities and estimating probabilities with the following transformation:

$$p = \frac{e^u}{1 + e^u}$$

where,
    p = probability of purchase
    e = the constant e
    u = product utility

7.4   **Randomized First Choice**

The Randomized First Choice (RFC) method combines many of the desirable elements of the First Choice and Share of Preference models. As the name implies, the method is based on the First Choice rule, and can be made to be immune to IIA difficulties. As with the Share of Preference model, the overall scaling (flatness or steepness) of the shares of preference can be tuned.

Most of the theory and mathematics behind the RFC model are nothing new. However, to the best of our knowledge, those principles have never been synthesized into a generalized conjoint/choice market simulation model. RFC, suggested by Orme (1998) and later refined by Huber, Orme and Miller (1999), was shown to outperform all other Sawtooth Software simulation models in predicting holdout choice shares for a data set they examined. The holdout choice sets for that study were designed specifically to include product concepts that differed greatly in terms of similarity within each set.

Rather than use the utilities as point estimates of preference, RFC recognizes that there is some degree of error around these points. The RFC model adds unique random error (variability) to the utilities and computes shares of preference in the same manner as the First Choice method. Each respondent is sampled many times to stabilize the share estimates. The RFC model results in a correction for product similarity due to correlated sums of errors among products defined on many of the same attributes. To illustrate RFC and how correlated errors added to product utilities can adjust for product similarity, consider the following example:

Assume two products: A and B. Further assume that A and B are unique. Consider the following product utilities for a given respondent:

	Avg. Product Utilities
A	10
B	30

If we conduct a first choice simulation, product B captures 100% of the share:

	Avg. Product Utilities	Share of Choice
A	10	0%
B	30	100%

However, let's assume that random forces can come to bear on the decision for this respondent. Perhaps he is in a hurry one day and doesn't take the time to make the decision that optimizes his utility. Or, perhaps product B is temporarily out-of-stock. Many random factors in the real world can keep our respondent from always choosing B.

We can simulate those random forces by adding random values to A and B. If we choose large enough random numbers so that it becomes possible for A to be sometimes chosen over B, and simulate this respondent's choice a great many times (choosing new random numbers for each choice iteration), we might observe a distribution of choices as follows:

	Avg. Product Utilities	Share of Choice
A	10	25.0%
B	30	75.0%

(Note: the simulation results in this section are for illustration, to provide an intuitive example of RFC modeling. We assume shares of preference are proportional to product utilities.)

Next, assume that we add a new product to the mix (A'), identical in every way to A. We again add random variability to the product utilities so that it is possible for A and A' to be sometimes chosen over B, given repeated simulations of product choice for our given respondent. We might observe shares of preference for the three-product scenario as follows:

	Avg. Product Utilities	Share of Choice
A	10	20.0%
A'	10	20.0% (A + A' = 40.0%)
B	30	60.0%

Because unique (uncorrelated) random values are added to each product, A and A' have a much greater chance of being preferred to B than either one alone would have had. (When a low random error value is added to A, A' often compensates with a high random error value). As a simple analogy, you are more likely to win the lottery with two tickets than with one.

Given what we know about consumer behavior, it doesn't make sense that A alone captures 25.0% of the market, but that adding an identical product to the competitive scenario should increase the net share for A and A' from 25.0% to 40.0% (the classic Red Bus/Blue Bus problem). It doesn't seem right that the identical products A and A' should compete as strongly with one another as with B.

If, rather than adding uncorrelated random error to A and A' within each choice iteration, we add the same (correlated) error term to both A and A', but add a unique (uncorrelated) error term to B, the shares computed under the first choice rule would be as follows:

	Avg. Product Utilities	Share of Choice
A	10	12.5%
A'	10	12.5% (A + A' = 25.0%)
B	30	75.0%

(We have randomly broken the ties between A and A' when accumulating shares of choice). Since the same random value is added to both A and A' in each repeated simulation of purchase choice, A and A' have less opportunity of being chosen over B as compared to the previous case when each received a unique error component (i.e. one lottery ticket vs. two). The final utility (utility estimate plus error) for A and A' is always identical within each repeated first choice simulation, and the inclusion of an identical copy of A therefore has no impact on the simulation result. The correlated error terms added to the product utilities have resulted in a correction for product similarity.

Let's assume that each of the products in this example was described by five attributes. Consider two new products (C and C') that are not identical, but are very similar—defined in the same way on four out of five attributes. If we add random variability to the part-worths (at the attribute level), four-fifths of the accumulated error between C and C' is the same, and only one-fifth is unique. Those two products in an RFC simulation model would compete very strongly against one another relative to other less similar products included in the same simulation. When C received a particularly large positive error term added to its utility, chances are very good that C' would also have received a large positive error term (since four-fifths of the error is identical) and large overall utility.

RFC Model Defined

We can add random variability at both the attribute *and* product level to simulate any similarity correction between the IIA model and a model that splits shares for identical products:

$$U_i = X_i (ß + E_{attribute}) + E_{product}$$

where:

U_i	=	Utility of product i for an individual or homogenous segment at a moment in time
X_i	=	Row of design matrix associated with product i
ß	=	Vector of part-worths

$E_{attribute}$ = Variability added to the part-worths (same for all products)
$E_{product}$ = Variability (Gumbel) added to product i (unique for each product)

Repeated draws are made to achieve stability in share estimates, computed under the First Choice rule. In RFC, the more variability added to the part-worths, the flatter the simulations become. The less variability added to part-worths, the more steep the simulations become. Under every possible amount of attribute variability (and no product variability), shares are split exactly for identical products, resulting in no "inflation" of net share. However, there may be many market scenarios in which some share inflation is justified for similar products. A second unique variability term (distributed as Gumbel) added to each product utility sum can tune the amount of share inflation, and also has an impact on the flatness or steepness of the overall share results. It can be shown that adding only product variability (distributed as Gumbel) within the RFC model is identical to the familiar logit model (Share of Preference Model). Therefore, any degree of scaling or pattern of correction for product similarity ranging between the First Choice model and Share of Preference can be specified with an RFC model by tuning the *relative* contribution of the attribute and product variability.

The exponent also can play a role in RFC, similar, but not identical to, product variability. Decreasing the exponent (multiplying the utility estimates by a value less than unity) decreases the variance of the utility estimates relative to the variance of the random variation added within RFC simulations, in turn making simulated shares flatter. There is a subtle difference between increasing product variability and lowering the exponent, though both result in a flattening of shares. If only attribute variation is being used in an RFC simulation, decreasing the exponent flattens the shares, but the overall model still does not reflect the IIA property. Adding product variability, however, flattens the shares and causes the RFC model to reflect at least some degree of IIA behavior. Though the exponent is not required to simulate different patterns of correction for product similarity and scaling, in the next section we show that it is useful to retain the exponent adjustment from an operational point of view.

The RFC model is very computationally intensive. With the suggested minimum of 100,000 total sampling iterations for a conjoint data set, the results tend to be fairly precise. But, if you have dozens of products in the simulation scenario, some product shares can become quite small, and greater precision would be needed. You can increase the precision by increasing the number of sampling iterations. The RFC model is appropriate for all types of conjoint simulations, based on either aggregate- or individual-level utilities. It provides the greatest benefits in the case of aggregate (logit and Lclass) models, which are more susceptible to IIA difficulties than individual-level models.

If you plan to use the RFC model with individual-level utilities to compute reasonably stable estimates of share at the *individual* level, you should sample each respondent at least 5,000 times, and preferably more. This can take a good deal of computing time. The software asks how many total sampling iterations to use, where the number refers to the total iterations across all respondents. If you have 500 respondents and want each respondent sampled 5,000 times, you should specify a total of 2,500,000 iterations.

The greatest complexity of the RFC model from an operational point of view is that the magnitude of the attribute variability multiplier must be adjusted whenever the number of products or number of attributes on which products differ changes across simulations to maintain comparable scaling of shares of preference. Our implementation of the RFC model includes an auto-calibrating attribute variability multiplier (developed through Monte-Carlo simulations) so that this issue is transparent to the user.

Ideally, you will provide your own validation data, in the form of holdout concepts or actual market shares, to permit calibration of the attribute and product variability terms. This allows you to best fit the scaling of shares and degree of correction for product similarity appropriate for the specific market you are modeling. If you are using the auto-calibrating attribute variability multiplier, you can adjust the relative contribution of attribute and product variability by manipulating only the

exponent and the product variability multiplier. For example, if you want to decrease the correction for product similarity by adding some product variability and adding less attribute-level variability to the utilities, you could increase the product variability multiplier (say, from 0.0 to 0.2). After doing so, the resulting shares of preference would be flatter than before the adjustment. (The more variation added to utilities, the flatter the resulting shares.) To "sharpen" the shares again, you would adjust the exponent upward. To adjust the amount of correction for similarity, you should have outside information about choices for holdout choice tasks, or actual market share, where the product concepts involved had significantly different degrees of product similarity within each set.

For more information about the RFC model, please refer to Huber, Orme and Miller's paper entitled, "Dealing with Product Similarity in Choice Simulations," available for downloading from our home page: http://www.sawtoothsoftware.com.

Index