

Chapter 10

Market Simulators for Conjoint Analysis

The market simulator is usually considered the most important tool resulting from a conjoint analysis 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 percentage of respondents projected to choose each product. 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. Market simulators are commercially available or can be constructed using spreadsheet programs.

10.1 What Is a Market Simulation?

A conjoint study leads to a set of utilities or part-worths that quantify respondents' preferences for each level of each attribute. These utilities can be analyzed in a number of ways. You can examine each respondent's utilities, but, if the number of respondents is large, this can be overwhelming. You might summarize the average utilities or compute average importances. You could create graphs and charts to display that information. But to many managers the results of conjoint analysis may seem abstract. Also, when we examine aggregate data or average responses, we may fail to detect important market segments—groups of consumers with 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 within competitive scenarios. The product concepts are defined in terms of the attributes and levels you used in the conjoint study. You walk into a virtual room, show them a market scenario (i.e., products A, B, and C), and they vote for the products 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 us suppose we were able to quantify how much people liked various flavors of ice cream. Let us refer to those preferences as utilities, and assume the following values for a given respondent:

<i>Flavor</i>	<i>Utility</i>	<i>Price</i>	<i>Utility</i>
Chocolate	0	\$0.60	50
Vanilla	30	\$0.80	25
Strawberry	40	\$1.00	0

Using these utility values, we can predict how this respondent would choose between a vanilla cone for \$0.80 or a strawberry cone for \$1.00:

$$\text{Vanilla (30 utiles) + \$0.80 (25 utiles) = 55 utiles}$$

$$\text{Strawberry (40 utiles) + \$1.00 (0 utiles) = 40 utiles}$$

We predict that this respondent will prefer the vanilla cone.

Now suppose we had data for not just one, but 500 respondents. We could count the number of times each of the two types of cones was preferred, and compute a share of preference, also referred to as a share of choice. If 300 respondents choose the vanilla cone for \$0.80 and 200 respondents choose the strawberry cone for \$1.00, then we would obtain these shares of preference or choice:

<i>Product Concept</i>	<i>Share of Choice</i>
Vanilla at \$0.80	$\frac{300}{500} = 0.60$
Strawberry at \$1.00	$\frac{200}{500} = 0.40$

The simplest market simulation is a simulation that assumes a first-choice model. A first-choice model assumes respondents buy or choose the product alternative from the competitive set that has the highest total utility, as determined by summing the part-worth utilities associated with the levels describing each product. There are more sophisticated approaches for market simulations that are beyond the scope of this introductory chapter. These more advanced approaches include logit, Bradley-Terry-Luce, and randomized first-choice models.

10.2 Applications of Conjoint Simulations

Looking only at average preferences or 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 targetable segments or the idiosyncratic behavior of individuals.

Consider the following example with three respondents and their preferences or utilities for color:

<i>Respondent</i>	<i>Blue</i>	<i>Red</i>	<i>Yellow</i>
Manny	50	40	10
Moe	0	65	75
Jack	40	30	20
Average	30	45	35

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

We can use simulators to answer basic questions about preference and shares of choice. We can use them to study the competitive environment and market segments. Furthermore, we can use the results of simulations to guide strategic decision making. Here are some of the benefits and applications of conjoint simulators:

- 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 almost every marketer cares about.
- As demonstrated earlier, conjoint simulations can capture idiosyncratic preferences occurring at the individual or group level. These underlying 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.
- Conjoint simulations can reveal differential substitutability (cannibalism or 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 utilities cannot reveal these important relationships.

- Conjoint simulations can reflect 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.
- Conjoint simulators may be used to answer questions about new products and new product introductions. 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 with one another. Each product is defined using the attribute levels measured in the conjoint study (brands, colors, prices, speeds, warranties, 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.
- Conjoint simulators may be used to guide pricing strategy. What is the relative price sensitivity of different brands? If I raise my price by 10 percent, how will it affect my brand? How will it affect competitor's 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.
- Conjoint studies can help us to answer questions about product bundles and product portfolios. What portfolio of products can I offer to appeal to different market segments and maximize overall share? If you have segmentation information (such as demographics or firmographics), 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, you can increase overall share for your product line or occupy a niche that is not currently being served.

The next three sections of this chapter will provide more detailed examples of applications, focusing upon introducing new products, estimating demand curves and elasticities, and designing products to appeal to market segments. For the examples in these sections you should assume the following three attributes, each with three levels:

<i>Brand</i>	<i>Style</i>	<i>Price</i>
A	X	\$100
B	Y	\$150
C	Z	\$200

10.3 Introducing New Products

Let us assume that your company is interested in entering a market that currently consists of just two competitors. There are only three attributes that adequately describe the products and account for preference in the market: brand, style, and price. The two products are Mellow (Brand A, Style X, at \$100) and Mild (Brand B, Style Y, at \$200).

Your company has developed a new product called Middling that has Style Z. You think Middling 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:

<i>Product</i>	<i>Brand</i>	<i>Style</i>	<i>Price</i>
Mellow	A	X	\$100
Mild	B	Y	\$200

Suppose a market simulation leads to the following shares of preference:

<i>Product</i>	<i>Share of Preference</i>
Mellow	64.3
Mild	35.7

In this simulation, we see that 64.3 percent of respondents preferred Mellow and 35.7 percent preferred Mild. 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 percent.

Let us 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: Middling (Brand C, Style Z, \$150. You add another product to your simulation specifications:

<i>Product</i>	<i>Brand</i>	<i>Style</i>	<i>Price</i>
Mellow	A	X	\$100
Mild	B	Y	\$200
Middling	C	Z	\$150

Running the simulation again might lead to the following shares of preference:

<i>Product</i>	<i>Share of Preference</i>
Mellow	42.5
Mild	21.3
Middling	36.2

You note that Mellow is still the most preferred product, but that your product Middling is preferred to Mild. Like any market research statistics computed from samples, shares of preference are not estimated without error. It is common to estimate a confidence interval to get a feeling for the degree of uncertainty due to sampling and measurement error associated with a given share of preference. Let us assume that the standard error reported for Middling in 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 Mellow. To simulate this, you lower the price slightly for your brand. Many simulators include the ability to interpolate between levels (straight line interpolation), so you can investigate even the smallest of price changes. As a first step, you decide to lower the price to \$130 for Middling (while holding the specifications for Mellow and Mild constant). The new simulated shares are as follows:

<i>Product</i>	<i>Share of Preference</i>
Mellow	39.2
Mild	19.0
Middling	41.8

You have overshot the mark (Middling's share exceeds Mellow's share), so you try a slightly higher price than \$130 and run the simulation again. You make repeated attempts until Middling's and Mellow's shares are equal. Let us 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 Middling commands a $\$136 - \$100 = \$36$ premium over Mellow. Respondents are indifferent between Brand A and Style X at \$100 and Brand C and Style Z at \$136.

10.4 Estimating Demand Curves and Elasticities

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

<i>Product</i>	<i>Brand</i>	<i>Style</i>	<i>Price</i>
Mellow	A	X	\$100
Mild	B	Y	\$200
Middling	C	Z	\$150

The shares of preference for the products, as defined above, were as follows:

<i>Product</i>	<i>Share of Preference</i>
Mellow	42.5
Mild	21.3
Middling	36.2

Let us assume that we wanted to estimate a demand curve for your company's offering: Middling, in the context of the current competition and prices. We do this through sensitivity analysis. Recall that we measured three distinct levels of price: \$100, \$150, and \$200. Note that we have already computed the share of preference for Middling when it is offered at \$150 (36.2). To estimate the demand curve for Middling, we will need to conduct two additional simulations: a simulation with Middling at the lowest price (\$100), and a simulation with Middling at the highest price (\$200). For each of these simulations, we'll hold the Mellow and Mild product specifications constant.

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

<i>Product</i>	<i>Brand</i>	<i>Style</i>	<i>Price</i>
Mellow	A	X	\$100
Mild	B	Y	\$200
Middling	C	Z	\$100

After running another simulation, we may observe the following shares:

<i>Product</i>	<i>Share of Preference</i>
Mellow	33.9
Mild	15.6
Middling	50.5

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

<i>Product</i>	<i>Brand</i>	<i>Style</i>	<i>Price</i>
Mellow	A	X	\$100
Mild	B	Y	\$200
Middling	C	Z	\$200

We run the simulation again, and the following shares are reported:

<i>Product</i>	<i>Share of Preference</i>
Mellow	49.2
Mild	26.9
Middling	23.9

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

<i>Middling</i> Price	<i>Middling</i> Share of Preference
\$100	50.5
\$150	36.2
\$200	23.9

We have demonstrated how to estimate a demand curve for Middling, 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 or middle price. It is often interesting to plot these demand curves and look at the patterns of price sensitivity among 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 elasticity using a single value, referred to as the price elasticity of demand:

$$E = \frac{\text{percentage change in quantity demanded}}{\text{percentage change in price}}$$

If the brand or product follows the law of demand, as 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 Middling:

$$E = \frac{\frac{(q_2 - q_1)}{(q_1 + q_2)/2}}{\frac{(p_2 - p_1)}{(p_1 + p_2)/2}}$$

$$E = \frac{\frac{(23.9 - 50.5)}{(50.5 + 23.9)/2}}{\frac{(200 - 100)}{(100 + 200)/2}} = \frac{-0.715}{0.667} = -1.073$$

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 must be 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). Even if this is the case, the relative price sensitivities for brands are still useful.

10.5 Designing Products for 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 products 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 called 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 you have enabled your market simulator to select respondents for analysis by segment, this can further enhance the power of the tool. For example, let's assume that a cluster analysis revealed three relatively different segments for the hypothetical example we've been using.

By examining the part-worths and importances 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 part-worth utility preferences in exhibit 10.1.

Attribute Level	Segment 1 (n = 128)	Segment 2 (n = 283)	Segment 3 (n = 216)
Brand A	39	-51	-44
Brand B	5	39	-29
Brand C	-44	12	73
Style X	61	-52	-34
Style Y	-23	45	-9
Style Z	-38	7	43
\$100	56	55	50
\$150	7	2	6
\$200	-63	-57	-56

Exhibit 10.1. Part-worth utilities across segments

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 market scenario we had used previously to obtain shares of preference across segments. Note that the shares below 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.

<i>Product</i>	<i>Brand</i>	<i>Style</i>	<i>Price</i>
Mellow	A	X	\$100
Mild	B	Y	\$200
Middling	C	Z	\$150

<i>Product</i>	<i>Shares of Preference</i>			
	<i>Segment 1</i> (<i>n</i> = 128)	<i>Segment 2</i> (<i>n</i> = 283)	<i>Segment 3</i> (<i>n</i> = 216)	<i>Total</i> (<i>n</i> = 627)
Mellow	84.8	21.5	22.2	34.7
Mild	7.4	40.0	14.2	24.5
Middling	7.8	38.5	63.6	40.8

Let us assume your company produces Old Middling under Brand C with Style Z at \$150. Your total share of preference is 40.8 percent. 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. Mellow, the Brand A product, clearly dominates Segment 1, which is the smallest segment.

Let us assume that your company was interested in offering an additional product, call it New Middling. We could examine the table of part-worth preferences in exhibit 10.1 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. Brand A, offering Style X at a low price, has got this relatively small segment nearly wrapped up, and this segment does not seem very receptive to Brand C.

You next consider Segment 2, which seems to represent a better opportunity for your brand. It is a relatively large segment that prefers Mild under Brand B, but also seems receptive to the Brand C product, Old Middling. Note also that Segment 2 strongly prefers Style Y, but your company currently offers only Style Z. By offering a Style Y product, you might be able to convert some current Brand B 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 us simulate what happens if, in addition to your current product Old Middling (Brand C, Style Z, \$150), you offer another product, New Middling (Brand C, Style Y, \$200).

Product	Shares of Preference			Total (<i>n</i> = 627)
	Segment 1 (<i>n</i> = 128)	Segment 2 (<i>n</i> = 283)	Segment 3 (<i>n</i> = 216)	
Mellow	82.2	17.2	18.6	31.0
Mild	7.2	32.0	11.9	20.0
Old Middling	6.8	27.7	47.8	30.4
New Middling	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 $[(30.4 + 18.7)/40.8] - 1 = 20$ percent in preference.

For line extension simulations you conduct, the answer will likely not be so clear and the process not 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. There are commercial software programs available that use different algorithms to find optimal or near-optimal solutions, even when the search space is extremely large. These optimizers use a variety of search algorithms, including exhaustive search, hill-climbing procedures, and genetic algorithms. Genetic algorithm and other search plug-ins for Excel are available, allowing researchers to construct their own simulators with optimization. More information on simulations and optimization approaches is available within the monograph by Krieger, Green, and Wind (2005).

10.6 Simulation Methods and Sample Sizes

Part-worth utilities can be used within a choice simulator to predict preference for different product concepts in competitive scenarios. There are various simulation methods, including the simple first-choice (maximum utility rule) and the logit or Bradley-Terry-Luce model. First-choice simulations assume that each respondent can choose or vote for only one product and that one alternative captures 100 percent of the share for each respondent. Shares of preference under the first-choice rule are proportions.

In contrast, logit or Bradley-Terry-Luce models let respondents choose products in a probabilistic manner. Suppose there are three products in a market scenario. Representing a respondent's preferences with a probabilistic model might show choice probabilities (0.6, 0.3, 0.1), but the first-choice rule would represent the probabilities as (1, 0, 0). The probabilistic model captures more information from each respondent and yields more stable share estimates. The standard errors for share predictions from logit or Bradley-Terry-Luce simulations are always smaller than under the first-choice rule. Therefore, if you plan to use the first-choice model, you will need larger sample sizes to stabilize share-of-choice estimates relative to probabilistic simulation models.

10.7 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 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. Also, using external effects often changes the fundamental properties of the original simulation model, such as the price sensitivities and substitution rates among products (Orme and Johnson 2006).

10.8 Multi-Store Simulators

The assumption of equal distribution is often responsible for the greatest differences between actual market shares and simulated shares of preference. Fortunately, there is a correct and straightforward simulation method for this problem. A multi-store simulator provides an appropriate way to account for an unequal distribution of products across the market without changing the products' original price sensitivities or substitution rates (Orme and Johnson 2006).

A multi-store simulator allows the researcher to specify, in the simplest case, the percentage of the regions/stores that carry each product. Superior implementations specify which products are available within each region/store and how much volume each region/store accounts for. Respondents are then randomly selected (with probability proportional to store volume) to make simulated visits to multiple stores on each of hundreds or thousands of occasions and to make choices among available products. If the respondent locations are known, we assign respondents to visit the applicable regional stores, rather than using a random process of assigning respondents to stores. The multi-store simulator is not just a tool for adjusting simulated shares to reflect better the availability of products across the market (and, in turn, market shares), but it is also a tool that more directly accounts for substitution effects by recognizing which products compete directly with one another (because they tend to be offered within the same regions/stores).