Chapter 9

Interpreting the Results
of Conjoint Analysis

Conjoint analysis provides various outputs for analysis, including part-worth utilities, counts, importances, and shares-of-preference simulations. This chapter discusses these measures and gives guidelines for interpreting results and presenting findings to management.

Before focusing on conjoint data, it is useful to review some fundamentals for interpreting quantitative data. The discussion of the nature of measurement scales follows the classic discussion of Stevens (1946), which has been adopted by numerous social scientists and business researchers. For current definitions and discussion, one can refer to a book on business statistics (Lapin 1993).

9.1   Nature of Quantitative Data

There are four general types of quantitative data:

- Nominal data. Here 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.
- Ordinal data. These 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 strengths of hurricanes. A category 3 hurricane is stronger and more damaging than a category 2 hurricane. It is generally not appropriate to apply arithmetic operations to ordinal data. The difference in strength between a category 1 and a category 2 hurricane is not necessarily equal to the difference in strength between a category 2 and a category 3. Nor can we say that a category 2 is twice as strong as a category 1 hurricane.

I would like to express special thanks to Rich Johnson for his contributions to this chapter in the section entitled "Price Elasticity, Price Sensitivity, and Willingness to Pay."

- Interval data. These permit the simple operations of addition and subtraction. The rating scales so common to market research provide interval data. The Celsius scale is an example of an interval scale. 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.

- Ratio data. These 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.

9.2   Conjoint Utilities

Conjoint utilities or part-worths are scaled to an arbitrary additive constant within each attribute and are interval data. The arbitrary origin of the scaling within each attribute is due to the nature of the experimental designs in conjoint analysis and the dummy coding in the design matrix. We could add a constant to the part-worths for all levels of an attribute or to all attribute levels in the study, and it would not change our interpretation of the findings.

When using a specific kind of dummy coding called effects coding, utilities are scaled to sum to zero within each attribute. A plausible set of part-worth utilities for fuel efficiency measured in miles per gallon might look like this:

Fuel Efficiency	Utility
30 mpg	-1.0
40 mpg	0.0
50 mpg	1.0

30 mpg received a negative utility value, but this does not mean that 30 mpg was unattractive. In fact, 30 mpg may have been acceptable to all respondents. But, all else being equal, 40 mpg and 50 mpg are better. The utilities are scaled to sum to zero within each attribute, so 30 mpg must receive a negative utility value. Other kinds of dummy coding set the part-worth of one level within each attribute to zero and estimate the remaining levels as contrasts with respect to zero.

Whether we multiply all the part-worth utilities by a positive constant or add a constant to each level within a study, the interpretation is the same. Suppose we have two attributes with the following utilities:

Color	Utility		Brand	Utility
Blue	30		A	20
Red	20		B	40
Green	10		C	10

The increase in preference from Green to Blue (twenty points) is equal to the increase in preference between brand A and brand B (also twenty points). However, due to the arbitrary origin within each attribute, we cannot directly compare values between attributes to say that Red (twenty utiles) is preferred equally to brand A (twenty utiles). And even though we are comparing utilities within the same attribute, we cannot say that Blue is three times as preferred as Green (30/10). Interval data do not support ratio operations.

9.3 Counts

When using choice-based conjoint (CBC), the researcher can analyze the data by counting the number of times an attribute level was chosen relative to the number of times it was available for choice. In the absence of prohibitions, counts pro-portions are closely related to conjoint utilities. If prohibitions were used, counts are biased. Counts are ratio data. Consider the following counts proportions:

Color	Proportion		Brand	Proportion
Blue	0.50		A	0.40
Red	0.30		B	0.50
Green	0.20		C	0.10

We can say that brand A was chosen four times as often as brand C (0.40/0.10). But, as with conjoint utilities, we cannot report that Brand A is preferred to Red.

9.4 Attribute Importance

Sometimes we want to characterize the relative importance of each attribute. We can 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 calculate percentages from relative ranges, obtaining a set of attribute importance values that add to 100 percent, as illustrated in exhibit 9.1. For this respondent, whose data are shown in the exhibit, the importance of brand is 26.7 percent, the importance of price is 60 percent, and the importance of color is 13.3 percent. 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.

Attribute	Level	Part-Worth Utility	Attribute Utility Range	Attribute Importance
Brand	A B C	30 60 20	60 - 20 = 40	(40/150) x 100% = 26.7%
Price	$50 $75 $100	90 50 0	90 - 0 = 90	(90/150) x 100% = 60.0%
Color	Red Pink	20 0	20 - 0 = 20	(20/150) x 100% = 13.3%

Utility Range Total
40 + 90 + 20 = 150

Exhibit 9.1. Relative importance of attributes

When summarizing attribute importances for groups, it is best to compute importances for respondents individually and then average them, rather than computing importances from average utilities. For example, suppose we were studying two brands, Coke and Pepsi. If half of the respondents preferred each brand, the average utilities for Coke and Pepsi would be tied, and the importance of brand would appear to be zero.

Importance measures are ratio-scaled, but they are also relative, study-specific measures. An attribute with an importance of twenty percent is twice as important as an attribute with an importance of ten, given the set of attributes and levels used in the study. That is to say, importance has a meaningful zero point, as do all percentages. But when we compute an attribute's importance, it is always relative to the other attributes being used in the study. And we can compare one attribute to another in terms of importance within a conjoint study but not across studies featuring different attribute lists.

When calculating importances from CBC data, it is advisable to use part-worth utilities resulting from latent class (with multiple segments) or, better yet, HB estimation, especially if there are attributes on which respondents disagree about preference order of the levels. (Recall the previous Coke versus Pepsi example.)

One of the problems with standard importance analysis is that it considers the extremes within an attribute, irrespective of whether the part-worth utilities follow rational preference order. The importance calculations capitalize on random error, and attributes with very little to no importance can be biased upward in importance. There will almost always be a difference between the part-worth utilities of the levels, even if it is due to random noise alone. For that reason, many analysts prefer to use sensitivity analysis in a market simulator to estimate the impact of attributes.

9.5 Sensitivity Analysis Using Market Simulations

Conjoint part-worths and importances may be difficult for nonresearchers to understand. Many presentations to management go awry when the focus of the conversation turns to explaining how part-worths are estimated or, given the scaling resulting from dummy variable coding, how one can or cannot interpret part-worths.

We suggest using market simulators to make the most of conjoint data and to communicate the results of conjoint analysis. When two or more products are specified in the market simulator, we can estimate the percentage of respondents who would prefer each. The results of market simulators are easy to interpret because they are scaled from zero to one hundred. And, unlike part-worth utilities, simulation results (shares of preference) are assumed to have ratio scale properties—it is legitimate to claim that a 40 percent share of preference is twice as much as a 20 percent share. Sensitivity analysis using market simulation offers a way to report preference scores for each level of each product attribute.

The sensitivity analysis approach can show us how much we can improve (or make worse) a product's overall preference by changing its attribute levels one at a time, while holding all other attributes constant at base case levels. We usually conduct sensitivity analyses for products assuming no reaction by the competition. In this way, the impact of each attribute level is estimated within the specific and appropriate context of the competitive landscape. For example, the value of offering a round versus a square widget depends on both the inherent desirability (utility) of round and square shapes and how many current competitors are offering round or square shapes. (Note that if no relevant competition exists or if levels needed to describe competitors are not included in the study, then it is possible to conduct sensitivity simulations considering the strength of a single product concept versus the option of purchasing nothing, or considering the product's strength in terms of purchase likelihood.)

Conducting sensitivity analysis starts by simulating shares of choice among products in a base case market. Then, we change product characteristics one level at a time (holding all other attributes constant at base case levels), We run the market simulation repeatedly to capture the incremental effect of each attribute level upon product choice. After we test all levels within a given attribute, we return that attribute to its base case level prior to testing another attribute.

To illustrate the method, we consider an example involving a study of mid-range televisions in 1997. The attributes in the study were as follows:

Brand

Sony
RCA
JVC

Screen Size

25-inch
26-inch
27-inch

Sound Capability

Mono Sound
Stereo Sound
Surround Sound

Channel Block Capability

None
Channel Blockout

Picture-in-Picture Capability

None
Picture-in-Picture

Price

$300
$350
$400
$450

Figure 9.1. Results of sensitivity analysis

Suppose we worked for Sony and the competitive landscape was represented by this base case scenario (3 products, each represented by rows of features in the table below):

Brand	Screen Size	Sound Capability	Channel Blockout Capability	Picture-in-Picture Capability	Price
Sony	25-inch	Surround	None	Picture-in-Picture	$400
RCA	27-inch	Stereo	None	Picture-in-Picture	$350
JVC	25-inch	Stereo	None	None	$300

Let us assume that, for this base case scenario, the Sony product captured 33 percent relative share of preference.

For a market simulation, we can modify the Sony product to have other levels of screen size, sound capability, channel blockout capability, and picture-in-picture capability, while holding the products from RCA and JVC constant. Figure 9.1 shows estimated shares of preference from this type of market simulation or sensitivity analysis. The potential (adjacent-level) improvements to Sony's product can be ranked as follows:

- Add channel blockout (48 relative preference)
- Reduce price to $350 (45 relative preference)
- Increase screen size to 26-inch (39 relative preference)

Sony cannot change its brand to RCA or JVC, so the brand attribute is irrelevant to management decision making (except to note that the Sony brand is preferred to RCA and JVC). And, although it is unlikely that Sony would want to reduce its features and capabilities, we can observe a loss in relative preference by including levels of inferior preference. One of those is price. Increasing the price to $450 results in a lower relative preference of 25 percent.

Before making recommendations to Sony management, we would, of course, conduct more sophisticated what-if analyses, varying more than one attribute at a time. Nonetheless, the one-attribute-at-a-time approach to sensitivity analysis provides a good way to assess relative preferences of product attributes.

9.6 Price Elasticity, Price Sensitivity, and Willingness to Pay

The results of conjoint analysis may be used to assess the price elasticity of products and services. It may also be used to assess buyer price sensitivity and willingness to pay. To begin this section, we should define price elasticity, price sensitivity, and willingness to pay:

- Price elasticity, by which we mean the price elasticity of demand, is the percentage change in quantity demanded divided by the percentage change in price. Price elasticity relates to the aggregate demand for a product and the shape of the demand curve. It is a characteristic of a product in a market.
- Price sensitivity is a characteristic of buyers or consumers. Some people are more sensitive to price changes than others, and the degree to which they are price sensitive can vary from one product or service to the next, one market to the next, or one time to the next. It can also vary with the characteristics of products described in terms of product attributes.
- Willingness to pay is a characteristic of buyers or consumers. A measure of willingness to pay shows how much value an individual consumer places on a good or service. It is measured in terms of money.

Conjoint analysis is often used to assess how buyers trade off product features with price. Researchers can test the price sensitivity of consumers to potential product configurations using simulation models based on conjoint results. Most often a simulation is done within a specific context of competitors. But when a product is new to the market and has no direct competitors, price sensitivity of consumers for that new product can be estimated compared to other options such as buying nothing.

The common forms of conjoint analysis measure contrasts between levels within attributes. The part-worths of levels are estimated on an interval scale with an arbitrary origin, so the absolute magnitudes of utilities for levels taken alone have no meaning. Each attribute's utilities are determined only to within an arbitrary additive constant, so a utility level from one attribute cannot be directly compared to another from a different attribute. To a trained conjoint analyst, an

array of utilities conveys a clear meaning. But that meaning is often difficult for an untrained audience to grasp. It is not surprising, then, that researchers look for ways to make conjoint utilities easier to interpret.

Monetary Scaling Trap

One common attempt to make conjoint utilities more understandable is to express them in monetary terms, or dollar equivalents. This is a way of removing the arbitrariness in their scaling. To do this, price must be included as an attribute in the conjoint design. Note that we cannot attach a monetary value to a single level (such as the color green), but must express the value in terms of differences between two colors, such as "green is worth $5 more than yellow." But if the attribute is binary (present/absent) such as "has sunroof" versus "doesn't have sunroof," the expressed difference is indeed the value of having the feature versus not having it.

The idea of converting utilities to dollar values can be appealing to managers. But some approaches to converting utilities to dollar equivalents are flawed. Even when computed reasonably, the results often seem to defy commonly held beliefs about prices and have limited strategic value for decision making.

Let us review a common technique for converting conjoint utilities to a monetary scale, and then we will suggest what we believe is a better approach. Here is how we can compute dollar equivalents from utilities. Imagine the following utilities for a single respondent for two attributes:

Attribute	Utility	Price	Utility
Feature X	2.0	$10	3.0
Feature Y	1.0	$15	1.0

For this respondent, a $5 change in price (from $15 to $10) reflects a utility difference of 2.0 (3.0 - 1.0). Therefore, every one utile change is equal to $2.50 in value (5 dollars/2.0 utiles). It then follows that feature X, being worth one utile more than feature Y, is also worth $2.50 more for this respondent.

We discourage the use of this type of analysis because it is a potentially misleading. Moreover, there is one practical problem that must be overcome if there are more than two price levels. Unless utility is linearly related to price, referencing different price points results in different measures of utiles per dollar. A common solution is to analyze the utility of price using a single coefficient. As long as the price relationship is approximately linear, this circumvents the issue.

Another problem arises when price coefficients are positive rather than negative as expected. This may happen for some respondents due to random noise in the data or respondents who are price insensitive. Such reversals would suggest willingness to pay more for less desirable features. One way to work around this is to compute dollar values of levels using average (across respondents) utilities, which rarely display reversals. Another approach to the problem is to ignore it,

assuming that the reversals are just due to random noise. A more proactive way to avoid reversals is to use an estimation method that enforces utility constraints, though there are potential drawbacks to this approach (Johnson 2000).

Additional complications arise when the price coefficient for a respondent is extremely small in absolute value, approaching zero. In that case, the dollar equivalents for incremental features become very large, approaching infinity. A typical way to handle this is to characterize the centers of the distributions using medians rather than means.

This type of analysis assumes that the conjoint method has accurately captured respondents' price sensitivity. Some conjoint methods (ACA and potentially any partial-profile method) tend to understate people's price sensitivity. This can result in inflated willingness to pay values.

But after taking the appropriate steps to compute reasonable dollar equivalents, the results are potentially misleading. Even when accurate price sensitivity has been estimated for each individual, an examination of average values will often reveal that respondents are willing to pay much more for one feature over another than is suggested by market prices. This often causes managers to disbelieve the results. However, we'll demonstrate later that such outcomes are to be expected when the monetary value of levels is computed in this way.

There are a number of fundamental problems with analysis based on average dollar values. First, it attempts to ascertain an average willingness to pay for the market as a whole. Firms usually offer products that appeal to specific targeted segments of the market. The firm is most interested in the willingness to pay among its current customers, or among buyers likely to switch to its products, rather than in an overall market average. Second, this approach does not reference any specific product, but instead considers an average product. We expect that a respondent's willingness to pay for an additional feature would depend upon the specific product that is being enhanced (e.g., a discount or a premium offering). Third, and most fundamental, this approach assumes no competition. Because a product purchase usually constitutes a choice among specific alternatives, the competitive context is a critical part of the purchase situation. To illustrate the fallacy of interpreting average dollar values, without respect to competitive offerings, consider the following illustration.

Economics on "Gilligan's Island"

Though perhaps loathe to admit it, many have watched the popular 1960s American TV program *Gilligan's Island*. The program revolved around an unlikely cast of characters who became marooned on an uncharted desert island. Nearly every episode saw the promise of rescue. And when it seemed that the cast was finally going to make it off the island, the bumbling Gilligan always figured out some way to ruin the day.

One colorful pair of characters were the ultrarich Mr. Howell and his wife. Now, imagine that one day a seaworthy boat with capacity for two passengers pulls into the lagoon and offers passage back to civilization for a price to be negotiated. What is the dollar value of rescue versus remaining on the island for Mr. and Mrs. Howell? Mr. Howell might pull out his checkbook and offer the crew millions of dollars. Under the assumption of no competition, the dollar equivalent utility of rescue is astronomically high. However, it might be much lower for other islanders of more limited means, and the average dollar value for all of them would have little relevance to the captain of the boat in negotiating a price. What would matter is the dollar value of the potential customers and no one else.

Now, assume, just as Mr. Howell and the first crew are preparing to shake on the deal, a second, equally seaworthy ship pulls into the lagoon and offers its services for a fixed $5,000. Ever the businessman, Mr. Howell will choose the $5,000 passage to freedom.

What has happened here? Is the utility of getting off the island for Mr. Howell suddenly different? Has his price sensitivity changed? No. The amount Mr. Howell would be projected to pay under the assumption of no competition is indeed very different from the amount he will pay given the appearance of another boat.

If the first boat's crew had administered a conjoint interview to Mr. Howell and had computed his willingness to pay under the first method reviewed in this chapter, they would have concluded that he was willing to pay a lot more than $5,000. But how meaningful is that information in light of the realities of competition? The realistic problem for the boat captain is to figure out what price the market will bear, given the existence of competitive offerings.

We can illustrate this point using another example. What is your willingness to pay for a color monitor for your laptop computer versus a monochrome screen? Assume we conducted a conjoint analysis including monochrome versus color monitors. If we computed your willingness to pay for color over monochrome, we would likely find that the incremental value of color over monochrome is worth a thousand dollars or more. But how meaningful is this information to a laptop manufacturer given the fact that laptops with color monitors are readily available on the market at quite inexpensive prices?

Price Sensitivity Simulations in Competitive Context

For most marketing problems involving competition, the best strategic information results from carefully defined market simulations. If a firm wants to assess the incremental demand resulting from offering specific features for its product, or improving its degree of performance, it should be estimated within a realistic competitive context.

Estimates of market demand should also be based on specific objectives. For example, the objective may be to determine how much more may be charged for a product or service by offering a new feature without any net loss in market acceptance. This approach involves simulating a realistic competitive scenario with a conjoint market simulator. Assume four products (A through D) represent the current relevant products in the marketplace. Further assume that the firm is interested in offering an additional feature for product A, and wants to estimate what new price can be charged while maintaining the same share of preference. We first simulate a base case with products A through D placed in competition with one another, where A does not include the new feature. We record its share of preference (say, 15 percent). We then conduct another simulation in which we improve A by offering a new feature (and hold the competition B through D constant). The share of preference for A should increase (say, to 20 percent). We then perform additional simulations (again holding competition constant) raising the price of the new product A until its share of preference again drops to the original 15 percent. The difference in price between the more expensive improved Product A that captures 15 percent and the old Product A that captured 15 percent reflects the incremental monetary value that the market will bear for the new feature, given the competitive context and the objective of maintaining share constant.

Market simulations conducted using individual-level utilities are best for this analysis. Individuals have different preferences, and the company that produces product A is most concerned with retaining current product A customers and attracting new buyers among those most likely to switch. The company does not care so much about individuals who are extremely unlikely to buy its offerings. Market simulations based on individual utilities support such complex market behavior, focusing the willingness-to-pay analysis on a relevant reference product and critical individuals rather than the market whole. Such market simulations can also reveal complex competitive relationships between products, such as degree of substitution (cross-effects) and differences in consumer price sensitivity to each product.

In summary, the common practice of converting differences between attribute levels to a monetary scale is potentially misleading. The value of product enhancements can be better assessed through competitive market simulations. If the market simulations are conducted using individual utilities, such simulations focus the price/benefit analysis on the customers that are most likely to purchase the firm's product(s) rather than on an overall market average. They provide strategic information based on a meaningful context that enables better decisions, while avoiding the pitfalls of other ways of analyzing data. Of course, the success of the simulation approach hinges on a number of assumptions, including the following: (1) the conjoint method produces accurate measures of price sensitivity, (2) the relevant attributes have been included in the simulation model, and (3) the relevant competitive offerings are reflected in the simulation model.