Pricing Tool Guide

Price Elasticity: Modeling Demand Functions

Once one has collected the pricing data, the next step is to build a model that fits the data connections between price and volume and can guide future price changes. There are options. One can build demand curves that do not make assumption about the underlying distribution of the data or one can assume that there is an underlying pattern. Because there are options, it is important to consider which one might not only be more accurate, but which might be more practical in terms of data collection and acceptance by those who must use it.

A curve that does not assume a constant pattern among the data could be one that, for instance, is based upon judgment. It may look "bumpy" after connecting the data points. This curve could take into account, for instance, that price could vary depending upon the price point. For instance, there might be certain threshold price points where price responses change as they are crossed. Ultimately, this form of demand curve may what one ends up with after considering the other, more function-based models. Pulling together a demand curve that fits the collective wisdom of those closest to the market can be done with a demand curve tool such as included in the Toolkit. This template leverages wisdom-of-the-crowds techniques by averaging individual responses separately and then averaging them together to form the curve. This template also has several checks for cannibalization assumptions and simple competitive response assumptions.

If, however, one wants a price-response model grounded in more consistent statistical logic, there are other options. One source of model variety is the choice of the statistical model itself. The more common models assume that the data follow linear, power, or logarithmic functions. These will be discussed below. The examples below come from the excellent book on pricing, Segmentation, Revenue Management, and Pricing Analytics, by Tudor Bodea and Mark Fergusons. [1] Another source of variety is the types of variables, sales, price, promotional quality, competitive activity, marketing mix, etc. A third form of variety is whether or not the data is transformed from its raw form in order to fit the assumptions required of the model, such as taking the logarithm of sales and/or price data. Multivariate or marketing mix models are not explored here.

Linear Demand:

The simplest form of the price response curve is a simple regression line that fits the untransformed sales and price data. For formula for this model is:

$$d(p) = D + m * p$$

with d(p) = the demand at price p; D = the demand at price = 0 (free), and m = the change in volume units per change in price units that best fits the data. Using a standard least squares regression method, such as one finds in the *ln()* function in Excel, one can find a linear relationship across the data like below. (Bodea, p 158-160) In the example

---

[1] Bodea, Tudor and Ferguson, Mark, Segmentation, Revenue Management, and Pricing Analytics, Routledge, New York,, 2014.

below from Bodea's book, the slope of -4.3 reflects the best fit across the data, which means that for every change in price units, the units of volume will change by -4.3. One can also see from this graph that the price elasticity (the rate of change in volume/ rate of change in price at any given point) varies throughout the curve even though the units change is constant.

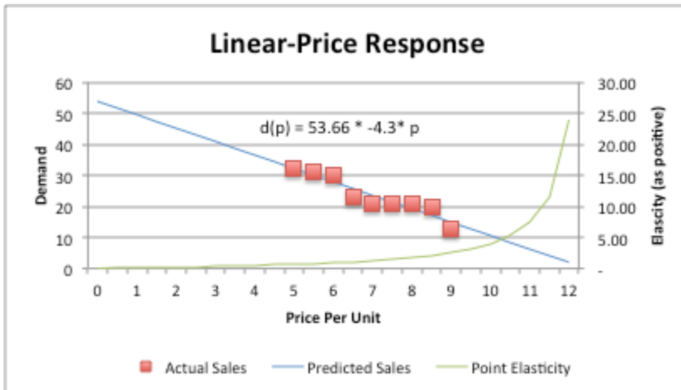

There are pros and cons with assuming linear relationships between price and volume. One of the main advantages is that it is simple to understand and apply, which is worth a lot. For this particularly date, there is also a good fit with an R^2 of .89. Often, however, there is not a good fit with linear assumptions like this. Another disadvantage is that the function also could be misinterpreted because it is unit dependent. For instance: the slope would be different if one measured volume change in gallons or liters. A more unit-less measure would be one based upon the rate of change in volume for a rate of change in price. This is the standard formula for price elasticity, which has some other uses such as checking for optimal pricing given one's costs. One sees from this graph that this elasticity measure varies throughout the model in a logarithmic progression.

Constant Elasticity Funtion:

The model that I have come across the most in my work is a multiplicative model that assumes constant elasticity across the data.

$$d(p) = C * p^{\varepsilon}$$

(c>0, ε <0)

With d(p) equals the demand at price (p), C = the demand at price = 0(free), p = the price, and e = the rate of change in price divided by the rate of change in volume. The example put forward in Bodea's book is from an in-store price experiment that yields the following graph.

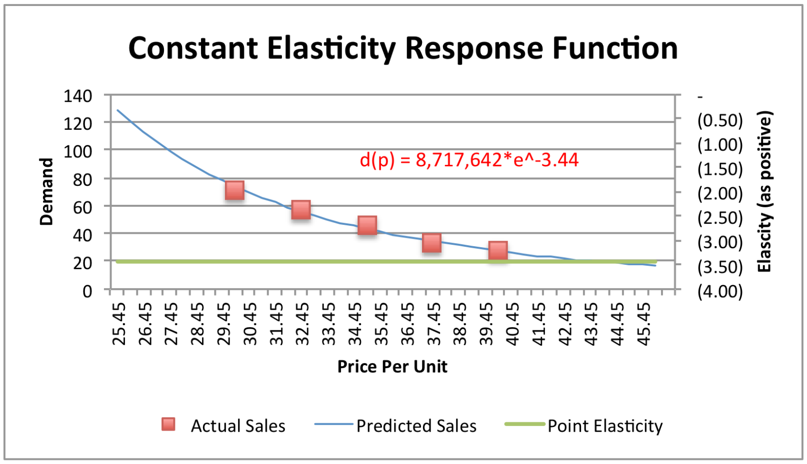

This model also has a number of pros and cons. While it appears intimidating, over small changes from the current price, **ε** can be treated as percent change in price vs. a percentage change in volume and can applied as a simple algebraic equation that is broadly interpretable. The constant price elasticity number also lends itself to other useful calculations, such as price and margin optimization checks. A disadvantage of this approach is that as one moves away from the current price, the effect in units for this formula moves exponentially towards the extremes. So, while it never is a good idea to extrapolate beyond the sample data, in my experience, this it if often done, and one drastically over or under estimate the impact of changes beyond small difference vs. from the current situation. One can see this effect in the example above.

Logarithmic Response Function:

A model that I have seen less commonly used is one that assumes a logarithmic relationship across the data. The formula for this is even more intimidating:

$$d(p) = C * \boldsymbol{\varepsilon}^{\frac{a+b*p}{q}} / (1 + \boldsymbol{\varepsilon}^{(\boldsymbol{a}+\boldsymbol{b}*\boldsymbol{p})})$$

(C >0, b <0)

Using this formula might have some appeal if the data, based upon a visualization of it, shows an S-curve demand relationship. It also has some intuitive appeal in that the elasticity flattens out at the extremes. An example of this formula from Bodea and Ferguson is shown below. Because of the need in this formula to solve for severable variable simultaneously, the example below was solved using an R program instead of just Excel to solve for the key variables. The ability to use more sophisticated software is another consideration when using this function. In the Bodea book, the examples shown also reflected discussions with management as to the right model to use. Interpretability was a key element in the ultimate choice. This seems like a best practice.

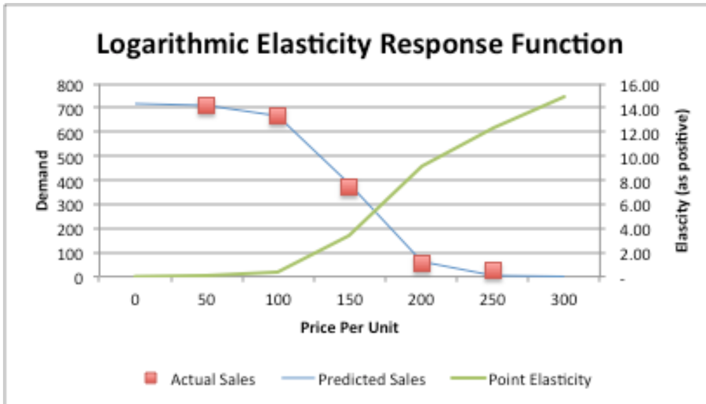

In conclusion, one can see that there are many ways to model the relationships between price and volume. This Guide only discussed some of the possibilities. As convenient as it might be to standardize on one model, the ideal is to be able to pick from several models in order to make the best inferences for future pricing decisions. At the end of the day, the final demand curve may very well include a hybrid of approaches. This brings one back to the importance of merging experience with the results of the data and analysis to provide the best advice on the impact of future changes. It also indicates why having someone skilled in modeling and the business context, who can guide the decision-makers through the decision process, is so valuable. It also indicates why this combination of skills is also not widely available.