Itron White Paper
Energy Forecasting

Price Elasticities in Energy Modeling and Forecasting

J. Stuart McMenamin, Ph.D.
Vice President, Itron Forecasting

Introduction

Economic theory and countless empirical studies support the idea that households and businesses respond to changes in the prices of goods and services. The response is implied by optimization behavior. In consumption decisions, individuals, in an effort to optimize their well being, will substitute away from goods that become more expensive toward goods that are relatively less expensive. Businesses, in an effort to minimize costs or maximize profits, will substitute away from more expensive inputs toward less expensive inputs. In either case, a price increase means a drop in the quantity demanded, and the strength of this response is often stated in terms of a price elasticity.

The goal of this paper is to provide background on price elasticity. The paper has the following sections:

- The End-Use Perspective, which provides a discussion of the market behavior and decision-making framework behind energy consumption decisions.
- Statistical Models and Functional Form, which discusses alternative model specifications and how elasticities are derived from model equations.
- Alternative Price Definitions, which covers the calculation of prices from multi-part utility tariffs and conversion of nominal prices to real prices.
- Review of Historical Electricity and Gas Prices, which shows the history of nominal and real prices over the last three decades.

The End-Use Perspective

The role of prices is somewhat complicated in the case of energy, because energy is an input into equipment that creates the final service being consumed. That is, we do not consume energy directly as we do a glass of milk. Instead, we consume a range of end-use services including lighting, heating, cooling, entertainment, communications, transportation and so on, all of which require energy inputs.

To understand the behavior that is behind short-term and long-term price response, it is necessary to think at the end-use level. The end-use framework suggests that there are five types of economic decisions where energy prices play a role:

- End-use acquisition decisions, which determine equipment saturation levels
- Fuel choice decisions in new construction, replacement and conversions
- End-use efficiency decisions at the time of equipment purchase
- Measure and device decisions that impact efficiency and usage
- Utilization levels

Of this list, only the final element—utilization—is variable in the short run. In the simplest case, for systems with manual on/off switches, an immediate response to higher prices is to reduce utilization. For example, for electric lighting, energy usage can be reduced directly by reducing the number of hours that lights are on. For more complex equipment, such as heating and cooling systems, utilization can be altered by changing settings on thermostats or other control devices, when such settings are available.

Equipment saturation levels and fuel shares change slowly over time, and contribute to the long-run elasticity. Saturation levels are typically determined by the value that an end-use service brings to the customer relative to the purchase and operating costs of the equipment. Often, the purchase cost of the equipment is the determining factor, since the energy costs will be relatively small. For example, few consumers would consider the energy costs of operating a computer or a telephone answering machine as part of the purchase decision.

Fuel-share decisions are important for a subset of uses, such as space heating, water heating, clothes dryers, cooking equipment and process heating equipment. In most cases, these decisions are long-run in nature, except in the case of dual-fuel equipment, which allows switching at a low cost once equipment is installed. Operating costs often play a very important role in the selection of heating and water heating systems as is evidenced in regional fuel usage patterns. In the Pacific Northwest, where electricity is relatively inexpensive, electric heating systems are prevalent. In the Northeast where electricity has been relatively expensive, fossil fuel systems are the dominant choice.

Efficiency decisions are also long-term in nature. These decisions are made at the time of equipment purchase, and it is rare to discard an inefficient piece of equipment when it still has a significant remaining life. In a world with low energy prices, equipment efficiency will not necessarily be a differentiating factor, but when operating costs become significant the market will provide a range of efficiency levels for major types of equipment. Government also plays a role here by establishing standardized efficiency measurement processes (such as mile per gallon ratings on cars) which allow consumers to make informed decisions.

When efficiency ranges are large, this can be a significant source of price elasticity. However, in recent years, the presence of government standards has limited the range of efficiencies available on the market for large appliances, such as refrigerators, water heaters, air conditioners and furnaces. By eliminating the inefficient options from the market, the consumer is faced with a restricted set of options. This reduces the long-term price elasticity to the extent that consumers are constrained by the standard.

Efficiency measures and devices include things like insulation levels, efforts to make buildings more air tight, set-back thermostats and occupancy sensors. While in most cases these items do not use energy directly, they do impact the amount of energy that is required to deliver a given level of service. Often utility programs target installation of efficiency measures as a way to help customers reduce consumption levels.

In the end-use modeling framework, the short- and long-run decisions are naturally segmented. Fuel and equipment decisions in new construction are modeled as a function of current or expected future fuel prices. Equipment conversion can also be modeled to account for customers who switch from fuel oil to natural gas or from a fossil fuel to a heat pump. Equipment efficiency is modeled as part of purchase and replacement decisions. And utilization is modeled as an ongoing process, based on prices, income, household size, and other economic and demographic factors.

Price elasticities in the end-use framework are embodied in the parameters of the equipment and utilization decision models. Typically, price elasticities are largest for the end-uses that are most directly controlled, such as heating and cooling, which can be altered through thermostat settings, and lighting, which can be reduced by turning lights off when not in use. Other end uses such as refrigerators, entertainment equipment, and communications are not typically modeled to have significant energy price elasticities.

Statistical Models and Functional Form

Detailed end-use models provide a useful framework for thinking about the types of behavior that determine price response. However, for most utilities, price elasticities are estimated using statistical models. These models relate historical consumption levels to economic variables, energy prices and weather. In general, these models can be represented as follows:

$$\text{EnergyUse} = F(\text{Economics}, \text{Weather}, \text{Energy Price}, \text{Other Factors}) + e \tag{1}$$

or more generally as:

$$Y = F(P, XOther) + e \tag{2}$$

where Y represents energy usage, P represents price, XOther is the set of remaining factors, and e is the statistical model error.

Definition of Price Elasticity. In this type of model, price response is determined by the derivative of energy use with respect to price. This derivative is sensitive to units of measurement. For example, if energy use is in MWh and price is in $/MWh, then the derivative has units of MWh per $/MWh. The elasticity normalizes the derivative to give a unitless measure of price response. The most general definition of a price elasticity is as follows:

$$Elast = \frac{dY}{dP} \times \frac{P}{Y} = \frac{dF(P, XOther)}{dP} \times \frac{P}{F(P, XOther) + e} \tag{3}$$

In this general form, it is clear that the elasticity is a function that must be evaluated for each level of P and for given values of the XOther variables. For a general nonlinear function, the elasticity value will change since changes in the derivative (dY/dP) and changes in the ratio (P/Y) will not cancel each other as price and other variables change.

Elasticity is often stated in terms of percentage changes. This flows naturally from equation (3), as seen below.

$$Elast = \frac{dY}{dP} \times \frac{P}{Y} = \frac{dY / Y}{dP / P} \approx \frac{\Delta Y / Y}{\Delta P / P} = \frac{\%\Delta Y}{\%\Delta P} \tag{4}$$

For very small percentage changes, this approximation is accurate, but for large percentage changes the result can be misleading.

Arc Elasticity. Another way to compute elasticities is using an “arc elasticity” formula. This formula takes two price/quantity points and gives the elasticity at the midpoint of a straight line between the two points on the demand curve. The arc elasticity is computed as:

$$ArcElast = \frac{F(P2) - F(P1)}{(P2 - P1)} \times \frac{(P2 + P1)/2}{(F(P2) + F(P1))/2} \tag{5}$$

The first part of this expression is a discrete version of the derivative. The numerator is the difference in the predicted values at two prices (deltaY) and the denominator is the difference in the price variable (deltaX). The ratio on the right is the average of the P values divided by the average of the Y values. If the demand relationship is a straight line between the two points, the arc elasticity is the point elasticity at the mid point of the demand segment.

Linear and Double Log Models. The most frequently used models are linear and double log models. The linear model can be represented as follows:

$$Y_t = a + b_p \times Price_t + \sum_i b_i \times XOther_{i,t} + e_t \tag{6}$$

With this form, the derivative (dY/dP) is the coefficient on p. As a result, the elasticity function can be written as:

$$\text{Elast} = \frac{dY}{dP} \times \frac{P}{Y} = b_p \times \frac{P}{Y} \tag{7}$$

An example is provided in Figure 1. In the figure, quantity demanded is depicted on the Y axis and price on the X axis. We expect the slope to be negative, so that as price increases, the quantity demanded decreases. (For those familiar with economics, the axes will seem to be reversed. In economics, price is typically placed on the Y axis. This is because the core of economics is the theory of value, in which demand and supply forces combine to determine a market price. For our purposes, however, it is simpler and less confusing if we put price (the causing factor) on the X axis and quantity (the outcome) on the Y axis.)

Several things are clear from the example. First, with a linear model, although the slope is constant (-.50 in the example), the elasticity is not. As prices rise and quantities decline, the price elasticity rises. Second, the arc elasticity, which is computed from two points on the demand curve, can be thought of as the elasticity at the midpoint of the linear segment between the two points.

The generalization to time series data is that the elasticity will be different in each month. Although it is common to compute the elasticity at the mean of the data (at average price and average quantity values), the elasticity of interest for forecasting is the end-of-period elasticity, which is computed based on conditions at the end of the historical data.

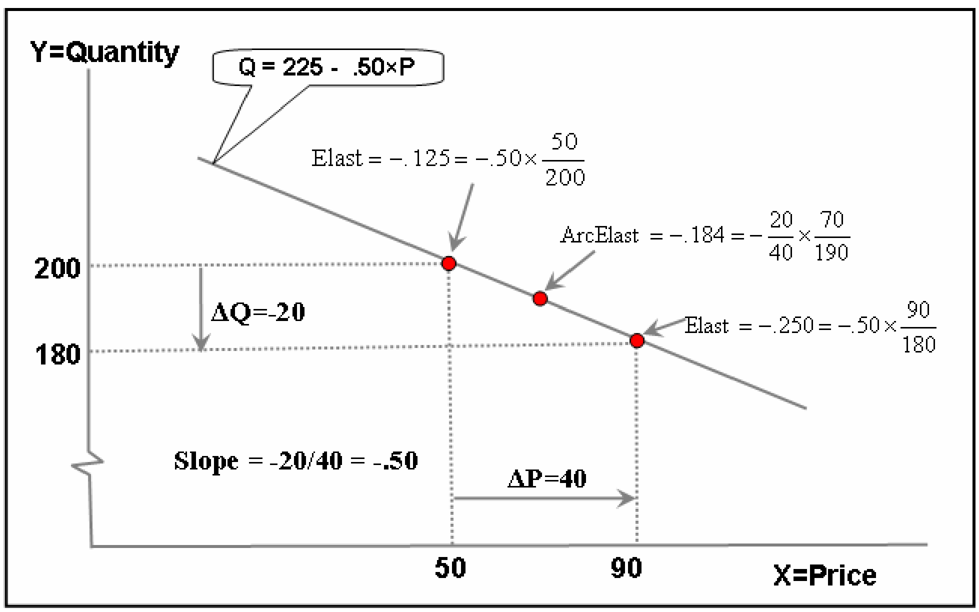

Figure 1 Price Elasticity for a Linear Equation

Double log equations are derived from geometric demand equations of the following form:

$$Y_t = a \times \text{Price}_t^{b_p} \times \prod_i \text{XOther}_t^{b_i} \times e_t \tag{8}$$

Taking logs of both sides gives the estimation form of the equation:

$$\ln(Y_t) = \ln(a) + b_p \times \ln(\text{Price}_t) + \sum_i b_i \times \ln(\text{XOther}_t) + \ln(e_t) \tag{9}$$

One of the properties of natural logs is that the derivative dln(X)/dX = 1/X. As a result, using the chain rule for derivatives, we get the following for the derivative of quantity with respect to price:

$$\frac{dY}{dP} = \frac{dY}{d\ln(Y)} \times \frac{d\ln(Y)}{d\ln(P)} \times \frac{d\ln(P)}{dP} = b_p \times \frac{Y}{P} \tag{10}$$

Substituting this slope formula into the elasticity definition gives:

$$\text{Elast} = \frac{dY}{dP} \times \frac{P}{Y} = b_p \tag{11}$$

As seen in (11), the elasticity for the double log form is a single constant number. With this specification, a one percent increase in price will result in a fixed percentage change in quantity, regardless of the location on the demand curve.

An example of a double log demand curve is shown in Figure 2. This curve is constructed to go through the same price/quantity points as the linear curve in Figure 1. The elasticity for this curve is -.179. As shown, the arc elasticity for the two demand points is unchanged, since this is the elasticity at the midpoint of the linear segment between the two points.

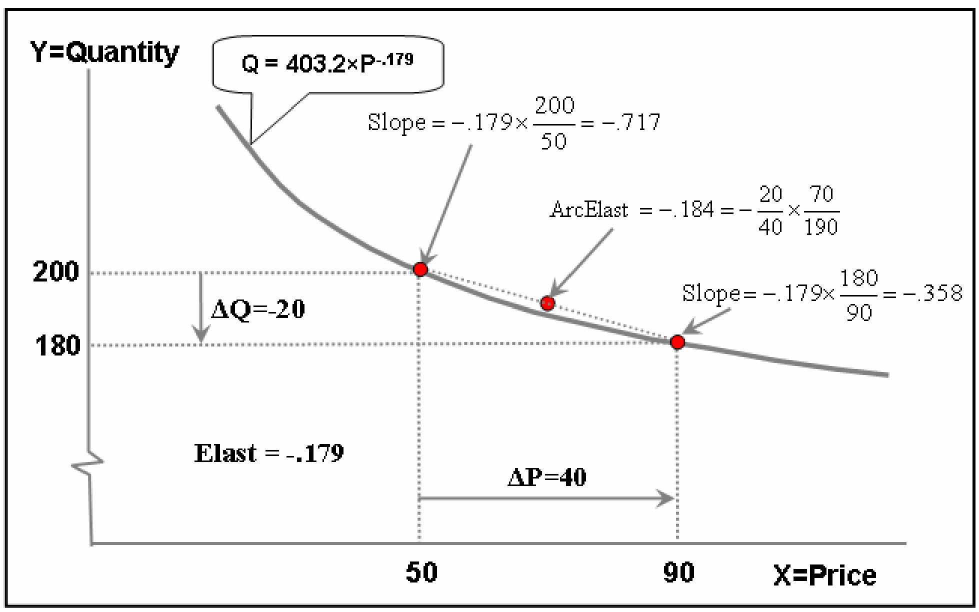

Figure 2 Price Elasticity for a Double Log Equation

Dynamic Specifications. So far, the price term is a single price variable with a time subscript. This is a static specification, in that the full price response is immediate. With a static specification, a shift in price has the same impact in later periods as it has in the first period impacted. To allow for dynamics in statistical models, it is necessary to introduce some form of distributed lag. The most common approaches are geometric or Koyck lags, use of moving averages, and polynomial distributed lags.

The introduction of geometric or Koyck lags can be based on stock adjustment models or geometric price expectation models. Either way, this translates into an estimated equation that includes a lagged dependent variable.

$$Y_t = a + b_p \times Price_t + \sum_i b_i \times XOther_{i,t} + c \times Y_{t-1} + e_t \qquad (12)$$

With this specification, the impact of a one-unit change in price is b in the first period. This feeds back in period 2 with an additional impact of $b \times c$. In period 3, there is an additional impact of $b \times c^2$. This continues with geometrically declining feedback effects, and the long run elasticity is:

$$ElastLR = b_p \times \left(1 + c + c^2 + c^3 + \ldots\right) \times \frac{P}{Y} = \frac{b_p}{1-c} \times \frac{P}{Y} \qquad (13)$$

With this specification, the long-run elasticity is a multiple of the short-run elasticity, and this multiple gets larger as the feedback parameter (c) increases. If c is .5, the long-run elasticity is twice the short-run elasticity. If c is .75, the long-run elasticity is four times the short-run elasticity.

One disadvantage of this specification is that all variables have the same lag structure. So variables that should have one period impacts (such as weather) also work through the same feedback coefficients. As a result, this approach is rarely used in weather-driven monthly energy equations.

As an alternative to an infinitely long geometric lag structure, the following specifies a finite lag structure of length L:

$$Y_t = a + b0_p \times Price_t + b1_p \times Price_{t-1} + bL_p \times Price_{t-L} + \sum_i b_i \times XOther_{i,t} + e_t \qquad (14)$$

From a practical perspective, this unconstrained or "free" lag structure will introduce more price parameters than can be supported by the estimation process. One useful simplification is to introduce additional structure through the use of moving averages.

$$Y_t = a + b0_p \times Price_t + b1_p \times \frac{\sum_{l=1}^{L} Price_{t-l}}{L} + \sum_i b_i \times XOther_{i,t} + e_t \qquad (15)$$

With this specification, the short-run elasticity results from the first price term. The long run price elasticity (after L periods) results from the sum of the two price coefficients. This type of specification is attractive, because price effects on equipment purchase decisions and on new construction decisions are spread across the lagged periods (e.g., prior ten years) with equal weight on each past value.

A final approach using polynomial distributed lags allows for less restrictive lag structures. With this approach, a functional form is selected for the lag structure (e.g., 3rd order polynomial with end constraints), and the parameters of the resulting polynomial are estimated.

$$Y_t = a + \sum_{l=1}^{L} PDL(l) \times Price_{t-l} + \sum_i b_i \times XOther_{i,t} + e_t \qquad (16)$$

This specification allows estimation of long distributed lags (e.g., 10 years is a 120 month lag with monthly data) with a relatively small number of coefficients (e.g., a 3rd order polynomial with a closing end constraint has three parameters).

The nature of price elasticity estimates also reflects the type of data that is used. With data for a single service territory, it is necessary to include dynamic terms to capture differences in short-run and long-run price response. With cross-sectional data, price elasticity estimates typically represent the long-run influence of differences in historical prices on historical equipment and construction decisions, as well as the influence of current prices on consumption behavior.

Alternative Price Definitions

Utility tariffs almost always include multiple parts. In the simplest case, there is a fixed charge and a usage charge. But more complex rates have energy blocks, demand charges, time-of-use energy charges and even real-time prices. Because of this complexity, it is necessary to decide how to define and compute price from historical revenue data or tariff elements.

From the customer perspective, information about price is communicated through monthly bills. Although these bills include information on the tariff elements, most residential and commercial customers do not look at or understand the details of the billing calculations. They do know when the bills go up or down, although for weather-driven loads, price changes may be masked by the seasonal variations. For electricity, the calculations are particularly abstract. It is not like gasoline, where we stand at the pump, see the price in $/gallon and have a firm understanding of what a gallon is and how far it will take us. Few people understand what a kilowatt hour is, and most people do not look at the price per kilowatt hour that is embodied in the total bill.

The three most commonly used measures of price in utility models are:

- Average Revenue. This is computed as revenue divided by volume.
- Marginal Price. The price of the last unit of energy for the typical consumer.
- Price Index. An index of tariff elements weighted into an index.

Average Revenue. This is the most commonly used measure of energy price. It is also the most problematic. Even in the case of simple rates, there is a built-in bias when average revenue is used because of the fixed charge. As consumption increases in summer or winter months, average prices fall as the fixed charge gets spread over a larger volume. So, volume increases cause average revenue (price) declines. This is true even if the underlying tariff elements remain fixed. In a statistical analysis, when we put average revenue on the right hand side of the equation, the direction of causality gets reversed, and it appears that the price declines caused the volume increases.

This problem is even worse in the case of declining block prices. An example of this situation is shown in Figure 3. In this example, there are three sets of underlying prices represented by the red diamonds, the blue circles and the green squares. Between rate cases, seasonal consumption patterns drive average revenue downward in the summer when cooling loads rise and downward in the winter as heating loads rise. Statistically, we want to sort out the impact of the two rate changes, but not be influenced by the seasonal variations within each rate regime.

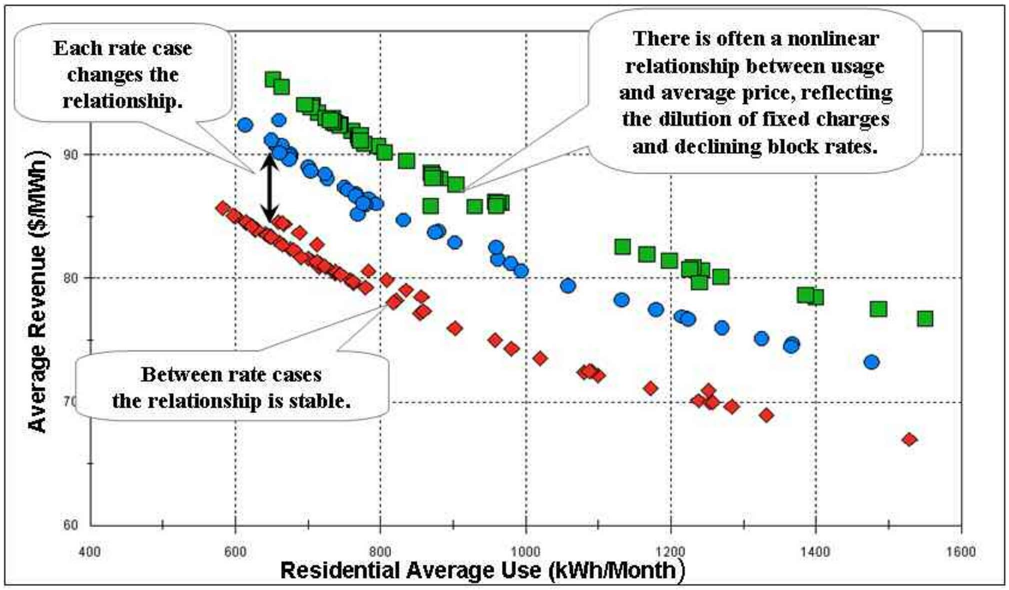

Figure 3 Average Revenue with Declining Blocks

The most common solution to this problem is to use a 12-month moving average of price as the price variable. This will remove most of the influence of seasonal price variations, resulting in three distinct price levels over the historical time period.

Marginal Price. The marginal price is the price of the last unit purchased. In economic theory, people do not make all or nothing decisions. They make decisions on the margin based on the value of the last unit consumed relative the cost of that last unit. Whether this thinking applies directly to energy consumption decisions is open to debate, but the use of marginal price eliminates the potential bias problem that occurs with average revenue. To implement this approach, a typical consumption level is chosen, and the marginal price or tariff at this consumption level is entered as the price variable. In the case in Figure 3, this would produce three distinct marginal price values with the change occurring when the underlying tariffs change.

Price Index. A price index is a way to derive a single price series from multiple components. One way to construct a price index for energy is to define a consumption pattern and price it out consistently over time. This could take the form of a typical utility bill. If the consumption pattern used to construct the typical bill has seasonal components, then it is important to take a 12-month moving average of the computed values to avoid the average revenue bias mentioned above.

Real vs. Nominal Prices. In time-series studies, it is important to convert price series to real terms. This is accomplished by dividing the energy price measure by an inflation index such as the Consumer Price Index or the GDP Deflator. When this is done, the typical profile will be that real prices decline gradually between rate cases, and increase when rates are adjusted upward. The same is true with economic variables, such as personal income and GDP, which should also be introduced in real terms.

Review of Historical Electricity and Gas Prices

Statistical estimation of price elasticities requires data on real energy prices and the corresponding volumes, as well as other factors that drive consumption. Elasticities are estimated from changes in prices and the corresponding changes in consumption. Estimation of significant elasticities requires significant historical price variations. The following two charts show 45 years of history for residential electric and gas prices. The data come from EIA and are constructed from electric and natural gas utility reports on residential revenues and residential sales volumes. The average price is constructed as the ratio of revenue to sales. Units are $ per MWh for electricity and $ per MCF for natural gas. The average prices in nominal terms are converted to real dollars by dividing by the consumer price index for all urban consumers (CPI-U), normalized to 1.0 in 2005.

Electricity Average Prices. As the electric chart shows, the nominal price of electricity was relatively flat between 1960 and 1973 at under $30 per MWh. There was a run-up of prices in the 1970s and early 1980s to a level above $70 per MWh. Since 1985, following the collapse of oil and gas prices, there has been a slow, steady rise to the current level of about $90 per MWh. Overall, the increase in average price from 1960 to 2005 is about 350 percent. Over the same period, the CPI has increased by 660 percent. As a result, real average residential electricity prices are currently at their lowest level in recorded history and have been stable at about $9.2 per MWh between 2000 and 2005.

Natural Gas Average Prices. The natural gas story is somewhat different. Starting at a little over $1 per MCF in the early 1960s, gas prices ramped up steeply through the early 1980s. Nominal and real prices peaked in the 1983-85 period. From 1983 to 1999, stable nominal prices resulted in declining real prices. Beginning in 2000, natural gas prices rose significantly to a level above $12 per MCF. As of 2005, in real terms, residential average prices were at their all-time high.

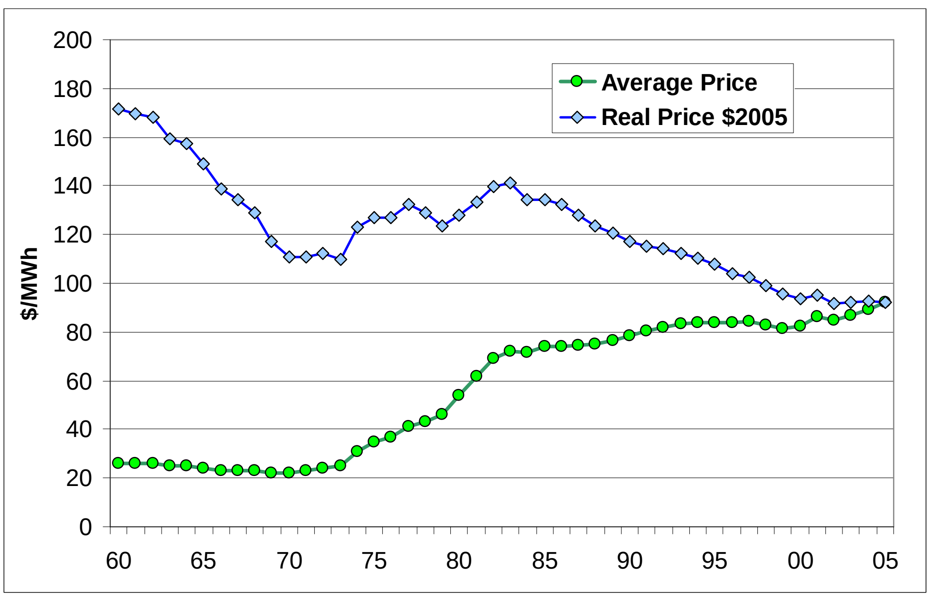

Figure 4 Real and Nominal Average Residential Electric Prices

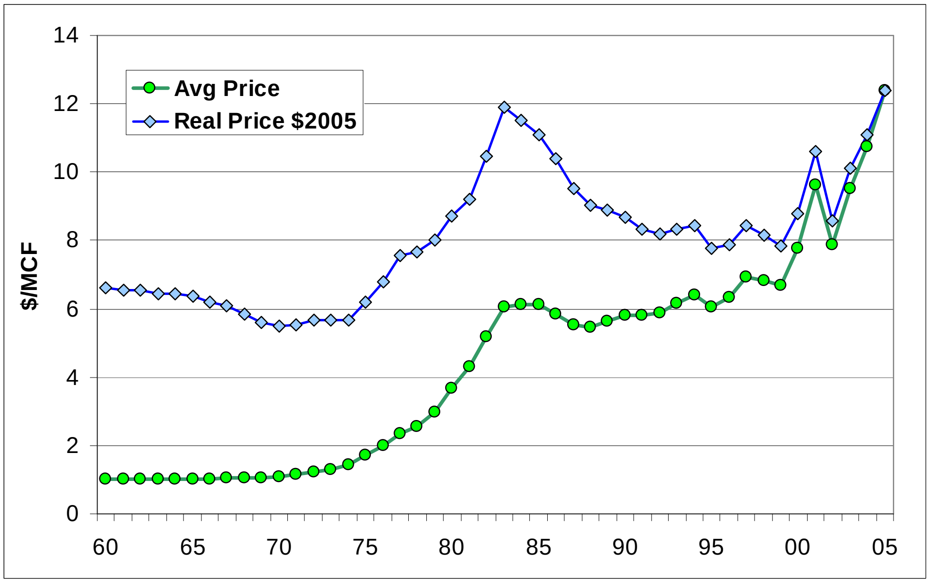

Figure 5 Real and Nominal Average Residential Gas Prices

For electricity, estimation of price elasticities over the recent years (since 1983) can be challenging. The time path of prices is a relatively steady decline in real terms. In absence of significant ups and downs, this series will be collinear with other modeling factors, such as a time trend or measures of average household income. As a result, estimation of price elasticities can be difficult for electricity.

For natural gas, recent history has significant swings in prices with a 33 percent decline in real prices between 1983 and the late 1990s, followed by a price spike in 2000 and a 50 percent real price increase over the last three years. As a result, statistical estimation of price elasticities is less challenging for natural gas.

Conclusion

Estimation of price elasticities and the use of these elasticities in energy forecasting was an important topic in the mid 1970s through the mid 1980s, a period during which real energy prices increased and during which there was significant uncertainty about the outlook for energy prices. Following 1985, the focus on price lessened as stable or declining prices led to an extended period of declining real energy prices. Fossil fuel price increases following 2000 have brought the issue of price elasticity back to the forefront, especially in natural gas markets. As a result, modeling of price response is again becoming an important component of energy forecasting. Reflecting this renewed interest, Itron completed a benchmarking study of utility practices early in 2006. The results of this study are summarized in a separate report which is available by request at forecasting@itron.com.

Itron Inc.

Itron is a leading technology provider and critical source of knowledge to the global energy and water industries. Nearly 3,000 utilities worldwide rely on Itron technology to deliver the knowledge they require to optimize the delivery and use of energy and water. Itron delivers value to its clients by providing industry-leading solutions for electricity metering; meter data collection; energy information management; demand response; load forecasting, analysis and consulting services; distribution system design and optimization; web-based workforce automation; and enterprise and residential energy management.

To know more, start here: www.itron.com

Itron Inc.
Corporate Headquarters
2111 North Molter Road
Liberty Lake, Washington 99019
U.S.A.
Tel.: 1.800.635.5461
Fax: 1.509.891.3355

Itron Inc.
Energy Forecasting - West
11236 El Camino Real
San Diego, California 92130-2650
Phone: 1.858.724.2620
Toll Free: 1.800.755.9585
Fax: 1.858.724.2690
Email: forecasting@itron.com
www.itron.com
www.itron.com/forecasting

Publication 100644WP-02
12/06