The Elasticity of Demand for Gasoline: A Semi-Parametric Analysis

Pin T. Ng
College of Business Administration
Northern Arizona University

James L. Smith
Cox School of Business
Southern Methodist University

Abstract

We use a semi-parametric conditional median as a robust alternative to the parametric conditional mean to estimate the gasoline demand function. Our approach protects against data and specification errors, and may yield a more reliable basis for public-policy decisions that depend on accurate estimates of gasoline demand. As a comparison, we also estimated the parametric translog conditional mean model. Our semi-parametric estimates imply that gasoline demand becomes more price elastic, but also less income elastic, as incomes rise. In addition, we find that demand appears to become more price elastic as prices increase in real terms.

KEY WORDS: Gasoline demand; Robust estimation; Quantiles; Splines; Semi-parametric model.

1 Introduction

Projections of future gasoline consumption are conditioned by the elasticity of demand. Thus, the design and success of various energy and environmental policy initiatives that pertain to gasoline necessarily involve judgments regarding this important aspect of consumer behavior. The magnitude of price elasticity, for example, largely determines the potency of excise taxes as a tool for raising government revenues, discouraging consumption and emissions, encouraging conservation and fuel switching, and attaining certain national security goals regarding energy independence. Moreover, the magnitude of income elasticity may influence the way in which economic growth and development affect progress towards achieving specific policy goals over time.

Numerous empirical studies have contributed to our understanding of the elasticity of demand for gasoline. Dahl and Sterner's (1991) very useful survey of previous research encompasses an almost bewildering variety of models, but finds a certain general consistency of results. They show, for example, that when appropriate allowances are made for differences in the treatment of dynamic adjustment processes and the effect of intervening variables, the preponderance of evidence suggests that gasoline demand is slightly inelastic with respect to price (the long-run elasticity being in the neighborhood of -0.9), but elastic with respect to income (approximately 1.2 in the long-run).

Although this body of previous research may suggest plausible consensus values that reflect the tendencies of representative consumers, the evidence is less conclusive on the question of whether it is appropriate to regard demand elasticities as being constant across groups of consumers who face different price and income levels. A study by McRae (1994), for example, shows that the demand for gasoline in the developing countries of Southeast Asia is somewhat less price elastic, but more income elastic, than that in the industrialized countries of the OECD. If this difference is due to variation in income levels between the two groups, we would then expect the rapid rate of increase in gasoline consumption that has been observed in the Asian countries to moderate as their incomes continue to rise. A more precise forecast, however, would require further knowledge of how elasticities vary with respect to price and income levels.

Most studies of gasoline consumption rely on models of the form:

$$Q = g(P, Y, Z) + \epsilon \tag{1}$$

which specifies the quantity of gasoline demanded Q as some unknown parametric function $g(\cdot)$ of the price of gasoline P, disposable income Y, and a vector of other explanatory variables Z, e.g., demographic characteristics, plus the disturbance term ϵ which captures the unexplained portion of demand. Economic theory provides information on the signs of partial derivatives of g, but not its functional form nor the specific nature of ϵ. Almost all analyses of gasoline demand to date, however, utilize some form of parametric specification on g (most notably linear, log-linear, or translog) and assume the distribution of ϵ to be normal with zero mean and fixed variance. The demand function g is then estimated by the conditional mean of Q using least-squares regression.

Although easy to apply, this method is not well suited for studying the potential variation in the elasticity of demand. The problem, of course, is that each functional specification "sees" a different pattern of variation in elasticities and imposes rigid constraints on what can be deduced from a given set of data. Reliance on the linear form forces estimated elasticities to vary hyperbolically, regardless of what the data might look like. Reliance on the log-linear form, on the other hand, is tantamount to assuming that elasticities are constant.

In this paper, we utilize a semi-parametric extension of the median regression technique of Koenker, Ng and Portnoy (1994) to study gasoline demand in the United States. This approach, which is based on tensor product polynomial splines, protects against misspecification in the demand functional form and achieves robustness against departures of the disturbance term from normality. Because we do not impose any predetermined structure on the demand function, the resulting estimates of demand elasticities, and their variation across price and income levels, reflect patterns of consumer choice that are inherent in the underlying data. We also develop and report confidence intervals that reflect the degree of uncertainty that is associated with the estimates that result from this semi-parametric procedure.

The only previous study that attempts a non-parametric specification of the gasoline demand function is by Goel and Morey (1993). They estimate the conditional mean using a kernel estimator and report considerable variation in the price elasticity of U.S. demand across the range of gasoline prices observed before and after the 1973 Arab oil embargo. However, they attribute all fluctuations in demand to variations in gasoline prices and ignore the influence of income and other variables. Therefore, the price elasticities which they report, which are sometimes *positive*, are probably contaminated by the confounding effects of omitted variables. Hausman and Newey (1992) also discuss a kernel estimator of the gasoline demand function, but since they do not report on elasticities, we cannot compare our results to theirs.

2 Theory and Data

Economic theory suggests a negative relationship between prices and quantity consumed. In addition, there is a positive income effect on consumption, at least if gasoline is a normal good. For the cross-sectional time-series data we use in this study, variations in population density across states also play an important role in determining consumption. For sparsely populated states like Montana, Nevada, New Mexico and Wyoming, where alternative forms of public transportation are not readily available, people rely heavily on the automobile as a means of transportation. The lack of close substitutes for automobile transportation suggests a relatively inelastic demand function.

The raw data spans 1952 to 1978 for forty-eight states. Alaska and Hawaii are dropped from the sample due to lack of data on gasoline prices. After 1978, gasoline was reclassified into regular, leaded, and unleaded grades, as well as full-service and self-service. Our sample, therefore, ends in 1978 to avoid inconsistent calibration in the data set. The gasoline prices and consumption used in this study are essentially those of Goel and Morey (1993). Quantity demanded is measured by gasoline consumption subject to taxation, as compiled by the *Federal Highway Administration*. Prices are the average service station gasoline prices within each state, as reported by *Platt's Oil Price Handbook and Oilmanac*. Annual per-capita personal income is taken from the *Survey of Current Business*. Population density is computed from state population divided by geographic areas, using figures from the *Statistical Abstract of the United States*. The price deflator is the consumer price index (1967 dollars) also from the *Statistical Abstract of the United States*.

Figure 1 presents scatter plots of annual gasoline consumption per capita in gallons (Q), prices per gallon in 1967 dollars (P), annual incomes per capita in 1967 dollars (Y) and population densities in thousand persons per square mile (D). To ameliorate the non-linearity of the data, we show a second set of scatter plots in Figure 2 with all variables measured in logarithmic form. In both figures, there appears to be a negative relationship between quantity and prices, a positive income effect, and a negative relationship between quantity and population density. There also seems to be a strong interaction between prices and income. The states with extremely sparse population (less than fifty persons per square mile) and high gasoline consumption (more than 750 gallons per person) are Montana, Nevada, New Mexico and Wyoming. In these states few means of alternative transportation are readily available. The scatter plot between Q and its one period lag Q_{-1} also suggests a significant inertia in gasoline consumption adjustment.

To illustrate the effect of per-capita income and population density on the demand curve, we present a series of conditional scatter plots of consumption per capita on prices over various intervals of income and density in Figure 3. The conditional plots allow us to see how quantity depends on prices given relatively constant values of the other variables (income and density). The different income ranges are given in the top panel while the given density levels are in the right panel in the figure. As we move from left to right, across a single row, income increases. Population densities rise as we move from the bottom to the top of each column. Also superimposed in the conditional scatter plots are cubic *basis spline* (B-spline) fits to the data in each panel [see de Boor (1978) or Schumaker (1981) for definition and construction of B-spline]. Figure 3 gives us a rough idea of how the demand curve looks over different income and population density regions. The conditional plots seem to suggest that both very low and very high income states have relatively price inelastic demand functions while the middle income and the moderately populated states have more elastic demand functions. Moving left to right horizontally across each row shows the positive income effect on consumption. Moving from bottom to top also reveals the negative population density effect on consumption. The highly populated states seem to have lower gasoline consumption per capita holding all else constant. We should, of course, emphasize that the conditional plots in Figure 3 only provide a very crude and tentative picture of the demand surface behavior.

3 Models and Estimation Techniques

3.1 The Semi-parametric Model

As we have seen in Figure 2, logarithmic transformations manage to remove a considerable amount of non-linearity from the data. Using lower case to denote logarithmically transformed variables, we adopt the following semi-parametric form of the demand equation:

$$q_i = \delta + g_{12}\left(p_i, y_i\right) + g_3\left(d_i\right) + \alpha q_{-1,i} + \epsilon_i \tag{2}$$

for $i = 1, \cdots, N$, where g_3 is a univariate continuous function, g_{12} is a bi-variate smooth function, α is a scalar and δ is the intercept term. The fact that g_3 and g_{12} are not confined to any specific parametric family provides flexibility to the model while the linear lag consumption component facilitates the traditional adaptive expectation dynamic analysis. The analysis of variance (ANOVA) decomposition in (2) has the virtue of separating the contributions of the covariates into main effects captured by the univariate functions and joint effect represented by the bi-variate function. It also provides a parsimonious representation of a potentially much more complex functional. Stone (1994) provides a theoretical justification for the use of ANOVA-type decomposition in multivariate function estimation using polynomial splines and their tensor products. We do not separate the main effect functions g_1 and g_2 for p and y from the joint-effect function g_{12} because the estimate of g_{12}, a tensor product bi-linear B-spline, spans the spaces of the univariate B-spline estimates $\hat{g}_1$ and $\hat{g}_2$. Including these main effect functions would introduce perfect multicollinearity and pose an identification problem during estimation. The population density d does not enter in interaction form because both economic theory and the data in Figure 1 seem to suggest that no interaction with other explanatory variables is warranted. Our parametric model chosen by the Akaike information criterion in Section 3.5 does not suggest any interaction between d and other explanatory variables either. We also estimated a fully nonparametric version of (2) with $\alpha q_{-1,i}$ replaced by a univariate continuous function $g_4(q_{-1,i})$. The result is identical to that of the semi-parametric specification. This supports our decision to model the lag of logarithmic consumption linearly.

3.2 Quantile Smoothing B-Splines

The demand equation (2) is estimated by solving the following optimization problem:

$$\min_{(\delta,\alpha)\in R^2; g_3, g_{12}\in\mathcal{G}} \sum_{i=1}^{N} \rho_\tau \left(q_i - \delta - g_{12}\left(p_i, y_i\right) - g_3\left(d_i\right) - \alpha q_{-1,i}\right) + \lambda \left\{ V_{12}\left(\frac{\partial g_{12}}{\partial p}\right) + V_{21}\left(\frac{\partial g_{12}}{\partial y}\right) + V_3\left(g_3'\right) \right\} \qquad (3)$$

where $\mathcal{G}$ is some properly chosen functional space, $\tau \in [0,1]$ specifies the desired conditional quantile, $\rho_\tau(u) = u(\tau - I(u < 0))$ is the check function which assigns a weight of τ to positive u and $\tau - 1$ otherwise, $\lambda \in (0, \infty)$ is the smoothing parameter, and V_3, V_{12} and V_{21} are measures of roughness to be defined below. For a given τ and λ, the estimated τth conditional quantile consumption function is

$$\hat{q}_{\tau,i} = \hat{\delta}_\tau + \hat{g}_{12_{\tau,\lambda}}\left(p_i, y_i\right) + \hat{g}_{3_{\tau,\lambda}}\left(d_i\right) + \hat{\alpha} q_{-1,i}$$

The special case of $\tau = .5$ yields the estimated conditional median consumption function. The portion in (3) associated with the check function controls the fidelity of the solution to the data while the roughness of the solution is measured by the total variations $V_3\left(g_3'\right) = \sum_{i=1}^{N-1} \left|g_3'\left(d_{i+1}\right) - g_3'\left(d_i\right)\right|$, over the mesh $M_d = \{di\}_{i=1}^{N}$ with knots $0 = d_0 < d_1 < \cdots < d_N < d_{N+1} = \infty$, and

$$V_{12}\left(\frac{\partial g_{12}}{\partial p}\right) = \sum_{i=1}^{N}\sum_{j=1}^{N-1}\left|\frac{\partial g_{12}}{\partial p}(p_{j+1}, y_i) - \frac{\partial g_{12}}{\partial p}(p_j, y_i)\right|$$

$$V_{21}\left(\frac{\partial g_{12}}{\partial y}\right) = \sum_{j=1}^{N}\sum_{i=1}^{N-1}\left|\frac{\partial g_{12}}{\partial y}(p_j, y_{i+1}) - \frac{\partial g_{12}}{\partial y}(p_j, y_i)\right|$$

over the tensor product mesh $M_p \otimes M_y$ constructed from the tensor products of the knots in $M_p = \{p_i\}_{i=1}^{N}$ and $M_y = \{y_i\}_{i=1}^{N}$. Notice that without loss of generality, we have assumed that d, p and y are already in ascending order to simplify notation. A discrete version of V_{12} and V_{21} appeared in a bivariate graduation paper of Portnoy (1994).

For an appropriately chosen $\mathcal{G}$, Koenker, Ng and Portnoy (1994) showed that, in the special case where there is only one covariate, the solution, which they called the τth quantile smoothing spline, is a linear smoothing spline, i.e. continuous piecewise linear function with potential breaks in the derivatives occurring at the knots of the mesh. With the linear smoothing spline characterization, the objective function (3) can be written in the form similar to that of the linear regression quantile in Koenker and Bassett (1978). This facilitates computation of the τth conditional quantile via modified versions of some familiar linear programs [See Koenker and Ng (1992) for a simplex algorithm and Koenker, Ng and Portnoy (1994) for a more general exposition]. Convergence rates of the quantile smoothing splines are given in He and Shi (1994), Portnoy (1994), and Shen and Wong (1994).

Even though computation of the quantile smoothing splines is feasible with an efficient linear program, it is still quite formidable for even a moderately large data set. In this paper, we suggest a B-spline approximation to the solution which utilizes a much smaller number of uniform knots in each mesh than those required in M_d and $M_p \otimes M_y$ hence saving tremendous memory and computing cycles. Details are given in the Appendix.

The fact that the solution is computed via a linear program leads to a very useful by-product — the entire family of unique conditional quantile estimates corresponding to the whole spectrum of τ can be computed efficiently by parametric programming [see Koenker, Ng and Portnoy (1994) for detail]. The same is true for the whole path of λ. This property will be exploited later in determining the optimal smoothing parameter and constructing the confidence interval of the conditional median estimate.

3.3 Choice of The Smoothing Parameter

The smoothing parameter λ in (3) balances the trade off between fidelity and roughness of the objective function. Its choice dictates the smoothness of the estimated conditional quantile. As $\lambda \to \infty$, the paramount objective is to minimize the roughness of the fit and the solution becomes the linear regression quantile of Koenker and Bassett (1978). On the other hand, when $\lambda \to 0$, we have a linear spline which interpolates every single q_i. We could, of course, assign different smoothing parameters to V_3 , V_{12} and V_{21} to produce different degrees of roughness along the direction of the separate covariates. Doing so, however, would complicate the choice of the correct smoothing parameters. In the single covariate problem, Koenker, Ng and Portnoy (1994) suggested using a modified version of Schwarz's (1978) information criterion (SIC) for choosing λ. The procedure is computationally feasible due to the univariate parametric programming nature of the problem in (3). Introducing more than one λ would require higher dimensional parametric programming.

The single smoothing parameter in the conditional median is chosen to minimize

$$SIC(\lambda) = \log\left(\frac{1}{N}\sum_{i=1}^{N}\left|q_i - \hat{\delta}_\tau - \hat{g}_{12_{\tau,\lambda}}(p_i, y_i) - \hat{g}_{3_{\tau,\lambda}}(d_i) - \hat{\alpha}_\tau q_{-1,i}\right|\right) + \frac{k(\lambda)}{2N}\log(N) \tag{4}$$

where k, which is inversely proportional to λ, is the effective dimensionality of the solution defined in Koenker, Ng and Portnoy (1994) [1]. The fidelity part of (4) can be interpreted as the log-likelihood function of the Laplace density. The second portion is the conventional dimension penalty for over-fitting a model.

The piecewise linear nature of the spline solution is particularly convenient for elasticity analysis. If the demand function is in fact log-linear, the optimal choice of λ will be very large and the demand function will be that characterized by the conventional log-linear model. If the demand function is only piecewise linear, the chosen λ will be relatively small and our quantile smoothing spline will produce a piecewise linear structure.

[1] The effective dimension of the estimated conditional median function is the number of interpolated q_i. Its value varies between N and the number of explanatory variables in (3) plus one (for the intercept). We can treat k as the equivalent number of independent variables needed in a fully parametric model to reproduce the semi-parametric estimated conditional median. When $k = N$, there is no degree of freedom and the estimated conditional median function passes through every response observation.

3.4 Confidence Set

Zhou and Portnoy (1994) suggested a *direct method* to construct confidence sets in the linear regression model

$$y_i = x_i'\beta + \epsilon_i \tag{5}$$

The $100\,(1-2\alpha)\,\%$ point-wise confidence interval for the τth conditional quantile at x' is given by

$$I_n = \left[x'\hat{\beta}_{\tau-b_n}, x'\hat{\beta}_{\tau+b_n}\right]$$

where $b_n = z_\alpha\sqrt{x'Q^{-1}x\tau\,(1-\tau)}$, $Q = \sum_{i=1}^n x_i x_i'$, z_a is the $(1-\alpha)$ quantile of the standard normal distribution, and $\hat{\beta}_{\tau-b_n}$ and $\hat{\beta}_{\tau+b_n}$ are the $(\tau - b_n)$th and $(\tau + b_n)$th regression quantiles of the linear regression model respectively. Utilizing the compact formulation in (9) in the appendix, we can adapt the direct method to compute the confidence sets of our estimated quantiles by treating the upper $N \times (K_d + K_p K_y + 2)$ partitioned matrix in (8) as the design matrix of (5). The lower $K \times (K_d + K_p K_y + 2)$ partition in (8) determines only the smoothness of the fitted quantile function and is irrelevant once an optimal λ has been chosen.

3.5 A Parametric Conditional Mean Model

To see how much our semi-parametric conditional median estimate differs from the conventional parametric conditional mean estimation, we fit the following translog model to the same data set:

$$q_{it} = \delta + \beta_1 p_{it} + \beta_2 y_{it} + \beta_{12} p_{it} y_{it} + \beta_3 d_{it} + \beta_4 d_{it}^2 + \alpha q_{it-1} + \epsilon_{it} \tag{6}$$

The model is chosen by minimizing the Akaike information criterion (AIC) in a stepwise model selection procedure.

4 Estimation Results

The SIC function of (4) for the whole spectrum of λ is plotted in Figure 4. The minimizing λ is .418 which occurs at an SIC of .021. The corresponding effective dimension k of the semi-parametric quantile smoothing B-splines fit is 17.

The quantile smoothing B-splines estimated α is .959 and hence the long-run elasticity multiplier is 24.4. As we have observed from the scatter plots in both Figure 1 and Figure 2, the inertia of short-run adjustment is quite high.

The dimension of the translog model (6) is 7, which is about 2/5 of the dimension ($k = 17$) of the semi-parametric model. This suggests that the semi-parametric model prefers a more complex structure than the translog model can offer. The least squares estimate of α is .95 and the long-run elasticity multiplier is 20 which is only slightly smaller than the semi-parametric estimate of 24.4.

We present the perspective plot of the estimated semi-parametric and translog demand surfaces conditioned at the median population density and lag consumption in Figures 5 and 18 respectively. There are altogether twenty-five grid points along both the price and income axes. The slightly positively sloped demand curves over the very low income region are the result of a boundary effect. There is just not enough data to make an accurate estimation near the boundary. As a result, in Figures 7 and 9, we discarded the first and last five grid points (the boundaries) and sliced through the demand surface at the different income grid points starting at the sixth and ending at the twentieth [2]. Also superimposed in the figures are the 95% point-wise confidence intervals. The boundary effect along the price direction is also reflected as wider intervals near the edges in each panel. The effect is more drastic for the semi-parametric model reflecting the slower convergence rate of the nonparametric approach.

The semi-parametric demand surface depicts the phenomenon we observed in the conditional plots in Figure 3. The price elasticity seems to be lower (in magnitude) when income is lower. As income increases, demand generally becomes more price sensitive. This confirms for U.S. consumers the type of income effect that McRae (1994) discovered in comparing price elasticities in industrialized and developing nations. Short-run price elasticity plots of the demand slices are shown in Figure 11. The price elasticity functions are step functions reflecting the piecewise linear nature of our quantile smoothing B-splines [3]. Short-run demand also seems to be less price elastic at lower prices and to become more elastic as prices increase. The translog model produces somewhat different result, as shown in Figure 13. As before, the price elasticity of demand appears to increase with income. However, the estimated price elasticity generally increases with price, which contradicts the semi-parametric estimates.

Our semi-parametric estimates of short-run price elasticity range (across price and income levels) from $-.205$ to $-.017$ with a median of $-.134$. This

[2]We plotted quantity on the vertical axis and price on the horizontal axis in Figures 7 and 9 to maintain coherence with Figures 5 and 18.

[3]Since we have measured all variables in logarithmic form, the slope of each log-linear segment of the demand curve corresponds to the elasticity.

compares to the average value of $-.24$ reported by Dahl and Sterner (1991) for models that use a comparable partial adjustment mechanism. Our short-run price elasticity evaluated at the median of the data is $-.205$, which is even closer to the number reported by Dahl and Sterner (1991). Our estimates of long-run price elasticity range from -5.02 to $-.415$, which extends well beyond the highest long-run price elasticities reported by Dahl and Sterner (1991). Our median estimate of -3.27 is about four times the magnitude of their average long-run elasticity.

Estimates of short-run price elasticity we obtain from the parametric translog model range from $-.295$ to $-.045$ with a median of $-.189$. The elasticity at the median of the data is $-.206$. In general, gasoline demand appears slightly more elastic when fitted to the translog model instead of the semi-parametric model.

A generally positive effect of income on consumption is also apparent in Figures 5 and 18. In Figures 15 and 19, we slice through the demand surface at fixed price levels to illustrate the Engel curves, and again superimpose the 95% confidence intervals in the plots. The slight negative income effects evident at extremely high price levels are artifacts of the data that can not be taken too seriously in view of the very wide confidence intervals that are found at the boundaries of the data set. The boundary areas are sparsely populated, as can be seen in all the four corners of Figure 24, which shows the projection of the data onto the price-income plane. The quantile smoothing B-splines estimates of short-run income elasticity corresponding to each panel in Figure 15 are presented in Figure 21 while estimates derived from the translog model are shown in Figure 23. Apart from the boundary effects noted above, short-run income elasticities seem to fall as income rises. This finding is also consistent with McRae's (1994) conclusion that income elasticities are generally lower in the relatively prosperous industrialized countries than in the developing countries of South East Asia.

Discarding the negative values near the boundary, our semi-parametric estimates of short-run income elasticity fall between 0 and .048, with a median value of .008. The short-run income elasticity evaluated at the median of the data is .042 which is substantially lower than the average short-run income elasticity of .45 reported by Dahl and Sterner (1991). Our estimates of long-run income elasticity range from 0 to 1.17, with a median value of .195. Our long-run income elasticity of 1.03 evaluated at the median of the data is fairly close to the average long-run elasticity of 1.31 which Dahl and Sterner (1991) report.

The estimates of short-run income elasticity we derive from the translog model range between 0 and .06 with a median of .01. These values are quite close to our semi-parametric estimates. They also confirm again the general tendency of income elasticities to decline as incomes rise. At the median of the data, the short-run income elasticity is .012, which is only 1/4 of that of the semi-parametric model.

5 Conclusion

We have applied a new, semi-parametric estimator to test the hypothesis that the elasticity of gasoline demand varies systematically across price and income levels. The approach we take uses the conditional median to produce a robust alternative to conventional parametric models that rely on the conditional mean. Our results tend to confirm, with different data and methods, both of McRae's (1994) suggestions: gasoline demand appears to become more price elastic, but also less income elastic, as incomes rise. In addition, we find that demand appears to become more price elastic as prices increase in real terms.

In comparison to previous parametric estimates of gasoline demand, our results tend to indicate that long-run adjustments are quite large relative to short-run effects. Regarding the effect of prices on demand, our short-run estimates are generally consistent with the consensus of previous short-run elasticities. Regarding the income effect, however, our long-run estimates tend to be much more consistent with the results of previous studies than are our short-run estimates. Further empirical research would be useful in helping to clarify what appears to be an important difference in the nature of dynamic adjustments to income versus price changes.

Acknowledgement

Computations in the paper were performed on computing facilities supported by the National Science Foundation Grants SES 89-22472. We thank Steve Craig for sharing his data on population and income, which was incorporated into the data set used in this study. We are also very grateful to Rajeev Goel for sending us his gasoline and consumption data. Any remaining errors in the data are, of course, our own responsibility. Finally, we thank Alok Bhargava and members of the University of Houston-Rice University Joint Econometric Seminar for many helpful comments on an earlier draft. Stimulating discussions with Roger Koenker and Steve Portnoy are also acknowledged.

Appendix: Quantile Smoothing B-splines

Let $\Delta_d = \{u_j\}_{j=1}^{K_d}$, $\Delta_p = \{v_j\}_{j=1}^{K_p}$ and $\Delta_y = \{w_j\}_{j=1}^{K_p}$ be uniform partitions of the domains covering d, p and y respectively. The B-spline representation of the estimate of g_3 is $\hat{g}_3(d_i) = \sum_{j=1}^{K_d} \hat{\gamma}_j B_j(d_i)$ and the bi-linear tensor product B-spline formulation of the estimate of g_{12} is

$$\hat{g}_{12}(p_i, y_i) = \sum_{k=1}^{K_p} \sum_{l=1}^{K_y} \hat{\gamma}_{k,l} B_k(p_i) B_l(y_i)$$

where

$$\gamma = \left(\gamma_{1,1}, \cdots, \gamma_{1,K_y}, \cdots, \gamma_{K_p,1}, \cdots, \gamma_{K_p,K_y}, \gamma_1, \cdots, \gamma_{K_d}\right)'$$

is the $(K_p K_y + K_d)$ vector of B-spline coefficients, and B_j, B_k, B_l are the linear basis functions defined over d, p, and y [see e.g. Schumaker (1981) for construction of the B-spline bases].

The objective function (3) can now be written as

$$\begin{aligned}
&\min_{\delta,\alpha,\gamma} \sum_{i=1}^{N} \rho\left(q_i - \delta - \sum_{k=1}^{K_p}\sum_{l=1}^{K_y} \gamma_{k,l} B_k(p_i) B_l(y_i) - \sum_{j=1}^{K_d} \gamma_j B_j(d_i) - \alpha q_{-1,i}\right) \\
&\quad + \lambda\left(\left[\sum_{i=1}^{K_y}\sum_{j=1}^{K_p-1}\left|\sum_{k=1}^{K_p}\sum_{l=1}^{K_y} \gamma_{k,l}\left\{B_k^{(1)}(v_{j+1}) - B_k^{(1)}(v_j)\right\} B_l(w_i)\right|\right]\right. \\
&\quad + \left[\sum_{j=1}^{K_p}\sum_{i=1}^{K_y-1}\left|\sum_{k=1}^{K_p}\sum_{l=1}^{K_y} \gamma_{k,l}\left\{B_l^{(1)}(w_{i+1}) - B_l^{(1)}(w_i)\right\} B_k(v_j)\right|\right] \\
&\quad + \left.\left[\sum_{i=1}^{K_d-1}\left|\sum_{j=1}^{K_d} \gamma_j\left\{B_j^{(1)}(u_{i+1}) - B_j^{(1)}(u_i)\right\}\right|\right]\right)
\end{aligned} \tag{7}$$

where $B_j^{(1)}$, $B_k^{(1)}$ and $B_l^{(1)}$ are first order derivatives of B_j, B_k and B_l respectively.

To express (7) as a linear program, we let $K = K_p(K_y - 1) + K_y(K_p - 1) + (K_d - 1)$, $\mathbf{1}$ be a vector of ones and $\mathbf{0}$ be a vector or matrix of zeros whose dimensions will be apparent from the context, $\theta' = (\delta, \gamma, \alpha)$ be a $(K_p K_y + K_d + 2)$ vector of parameters,

$$\mathbf{G}^{12} = \begin{bmatrix} \mathbf{G}_1^{12} \\ \vdots \\ \mathbf{G}_N^{12} \end{bmatrix}$$

be an $N \times K_p K_y$ matrix with rows

$$\begin{aligned}
\mathbf{G}_i^{12} \quad = \quad & \left(B_1(p_i) B_1(y_i), \cdots, B_1(p_i) B_{K_y}(y_i), \cdots, B_{K_p}(p_i) B_1(y_i), \cdots,\right. \\
& \left. B_{K_p}(p_i) B_{K_y}(y_i)\right)
\end{aligned}$$

for $i = 1, \cdots, N$,

$$\mathbf{G}^3 = \begin{bmatrix} B_1(d_1) & \cdots & B_{K_d}(d_1) \\ \vdots & \cdots & \vdots \\ B_1(d_N) & \cdots & B_{K_d}(d_N) \end{bmatrix}$$

be an $N \times K_d$ matrix,

$$\mathbf{V}^{12} = \begin{bmatrix} \mathbf{V}^{12}_{1,1} \\ \vdots \\ \mathbf{V}^{12}_{1,K_y} \\ \vdots \\ \mathbf{V}^{12}_{K_p-1,1} \\ \vdots \\ \mathbf{V}^{12}_{K_p-1,K_y} \end{bmatrix}$$

be a $(K_p - 1)(K_y) \times K_p K_y$ matrix in which

$$\begin{aligned} \mathbf{V}^{12}_{i,j} \quad = \quad & \left(\left\{ B_1^{(1)}(v_{i+1}) - B_1^{(1)}(v_i) \right\} B_1(w_j), \cdots, \right. \\ & \left\{ B_1^{(1)}(v_{i+1}) - B_1^{(1)}(v_i) \right\} B_{K_y}(w_j), \cdots, \\ & \left\{ B_{K_p}^{(1)}(v_{i+1}) - B_{K_p}^{(1)}(v_i) \right\} B_1(w_j), \cdots, \\ & \left. \left\{ B_{K_p}^{(1)}(v_{i+1}) B_{K_p}^{(1)}(v_i) \right\} B_{K_y}(w_j) \right) \end{aligned}$$

for $i = 1, \cdots, K_p - 1$ and $j = 1, \cdots, K_y$,

$$\mathbf{V}^{21} = \begin{bmatrix} \mathbf{V}^{21}_{1,1} \\ \vdots \\ \mathbf{V}^{21}_{K_p,1} \\ \vdots \\ \mathbf{V}^{21}_{1,K_y-1} \\ \vdots \\ \mathbf{V}^{21}_{K_p,K_y-1} \end{bmatrix}$$

be a $(K_p)(K_y-1)\times K_pK_y$ matrix in which

$$\begin{aligned}\mathbf{V}^{21}_{i,j} &= \Big(\Big\{B_1^{(1)}(w_{i+1})-B_1^{(1)}(w_i)\Big\}B_1(v_j),\cdots,\\ &\Big\{B_1^{(1)}(w_{i+1})-B_1^{(1)}(w_i)\Big\}B_{K_p}(v_j),\cdots,\\ &\Big\{B_{K_y}^{(1)}(w_{i+1})-B_{K_y}^{(1)}(w_i)\Big\}B_1(v_j),\cdots,\\ &\Big\{B_{K_y}^{(1)}(w_{i+1})\,B_{K_y}^{(1)}(w_i)\Big\}B_{K_p}(v_j)\Big)\end{aligned}$$

for $i=1,\cdots,K_p$ and $j=1,\cdots,K_y-1$,

$$\mathbf{V}^3=\begin{bmatrix}\Big\{B_1^{(1)}(u_{i+1})-B_1^{(1)}(u_i)\Big\} & \cdots & \Big\{B_{K_d}^{(1)}(u_{i+1})-B_{K_d}^{(1)}(u_i)\Big\}\\ \vdots & \cdots & \vdots\\ \Big\{B_1^{(1)}(u_{i+1})-B_1^{(1)}(u_i)\Big\} & & \Big\{B_{K_d}^{(1)}(u_{i+1})-B_{K_d}^{(1)}(u_i)\Big\}\end{bmatrix}$$

be a $(K_d-1)\times K_d$ matrix,

$$\mathbf{V}=\begin{bmatrix}\mathbf{V}^{12}\\ \mathbf{V}^{21}\end{bmatrix}$$

be a $(K-K_d+1)\times K_pK_y$ matrix,

$$\tilde{X}=\begin{bmatrix}\mathbf{1} & \mathbf{G}^{12} & \mathbf{G}^3 & \mathbf{1}\\ \mathbf{0} & \mathbf{V} & \mathbf{0} & \mathbf{0}\\ \mathbf{0} & \mathbf{0} & \mathbf{V}^3 & \mathbf{0}\end{bmatrix} \tag{8}$$

be an $(N+K)\times(K_pK_y+K_d+2)$ pseudo design matrix, $\tilde{y}=(q_1,\cdots,q_N,\mathbf{0})'$ be a $(N+K)$ pseudo response vector, and

$$\omega=\begin{bmatrix}\frac{1}{2}+\left(\tau-\frac{1}{2}\right)sgn\left(\tilde{y}_1-\tilde{x}_1\theta\right)\\ \vdots\\ \frac{1}{2}+\left(\tau-\frac{1}{2}\right)sgn\left(\tilde{y}_N-\tilde{x}_N\theta\right)\\ \lambda\mathbf{1}\end{bmatrix}$$

be an $(N+K)$ vector of weights. We can then express (7) in the following compact form:

$$\min_{\theta\in R^{K_pK_y+K_d+2}}\sum_{i=1}^{N+K}\omega_i\left|\tilde{y}_i-\tilde{x}_i\theta\right| \tag{9}$$

This is a variant of the minimization problem for the linear regression quantile of Koenker and Bassett (1978) and can be solved by the modified linear programming algorithm in Koenker and Ng (1992).

References

[1] Dahl, Carol and Thomas Sterner, (1991), Analysing gasoline demand elasticities: A survey, *Energy Economics*, July , 203-210.

[2] de Boor, Carl, (1978), *A Practical Guide to Splines*, Springer-Verlag, New York.

[3] Goel, Rajeev K. and Mathew J. Morey, (1993), Effect of the 1973 oil price embargo: A non-parametric analysis, *Energy Economics*, Jan., 39-48.

[4] Hausman, Jerry A. and Whitney K. Newey, (1992), Nonparametric estimation of exact consumers surplus and deadweight loss, preprint.

[5] He, Xuming and Peide Shi, (1994), Convergence rate of B-spline estimators of nonparametric conditional quantile function, *Journal of Nonparametric Statistics*, 3, 299-308.

[6] Koenker, Roger and Gilbert W. Bassett, (1978), Regression quantiles, *Econometrica*, 46, 33-50.

[7] Koenker, Roger and Pin T. Ng, (1992), Computing quantile smoothing splines, *Computing Science and Statistics*, 24, 385-388.

[8] Koenker, Roger, Pin Ng and Stephen Portnoy, (1994), Quantile smoothing splines, *Biometrika*, 81, 673-680.

[9] McRae, Robert, (1994), Gasoline demand in developing Asian countries, *The Energy Journal*, 15, 143-155.

[10] Portnoy, Esther, (1994), Bivariate Schuette graduation, *Actuarial Research Clearing House*, 1994.1, 127-134.

[11] Portnoy, Stephen, (1994), Local Asymptotics for quantile smoothing splines, preprint.

[12] Schumaker, Larry L., (1981), *Spline Functions: Basic Theory*, John Wiley & Sons, New York.

[13] Schwarz, G., (1978), Estimating the dimension of a model, *Annals of Statistics*, 461-464.

[14] Shen, Xiaotong and Wong, W. H., (1994), Convergence rate of sieve estimates, *Annals of Statistics*, (in press).

[15] Stone, Charles J., (1994), The use of polynomial splines and their tensor products in multivariate function estimation, *Annals of Statistics*, 22, 118-171.

[16] Zhou, Qing and Stephen L. Portnoy, (1994), Direct use of regression quantiles to construct confidence sets in linear models, preprint.

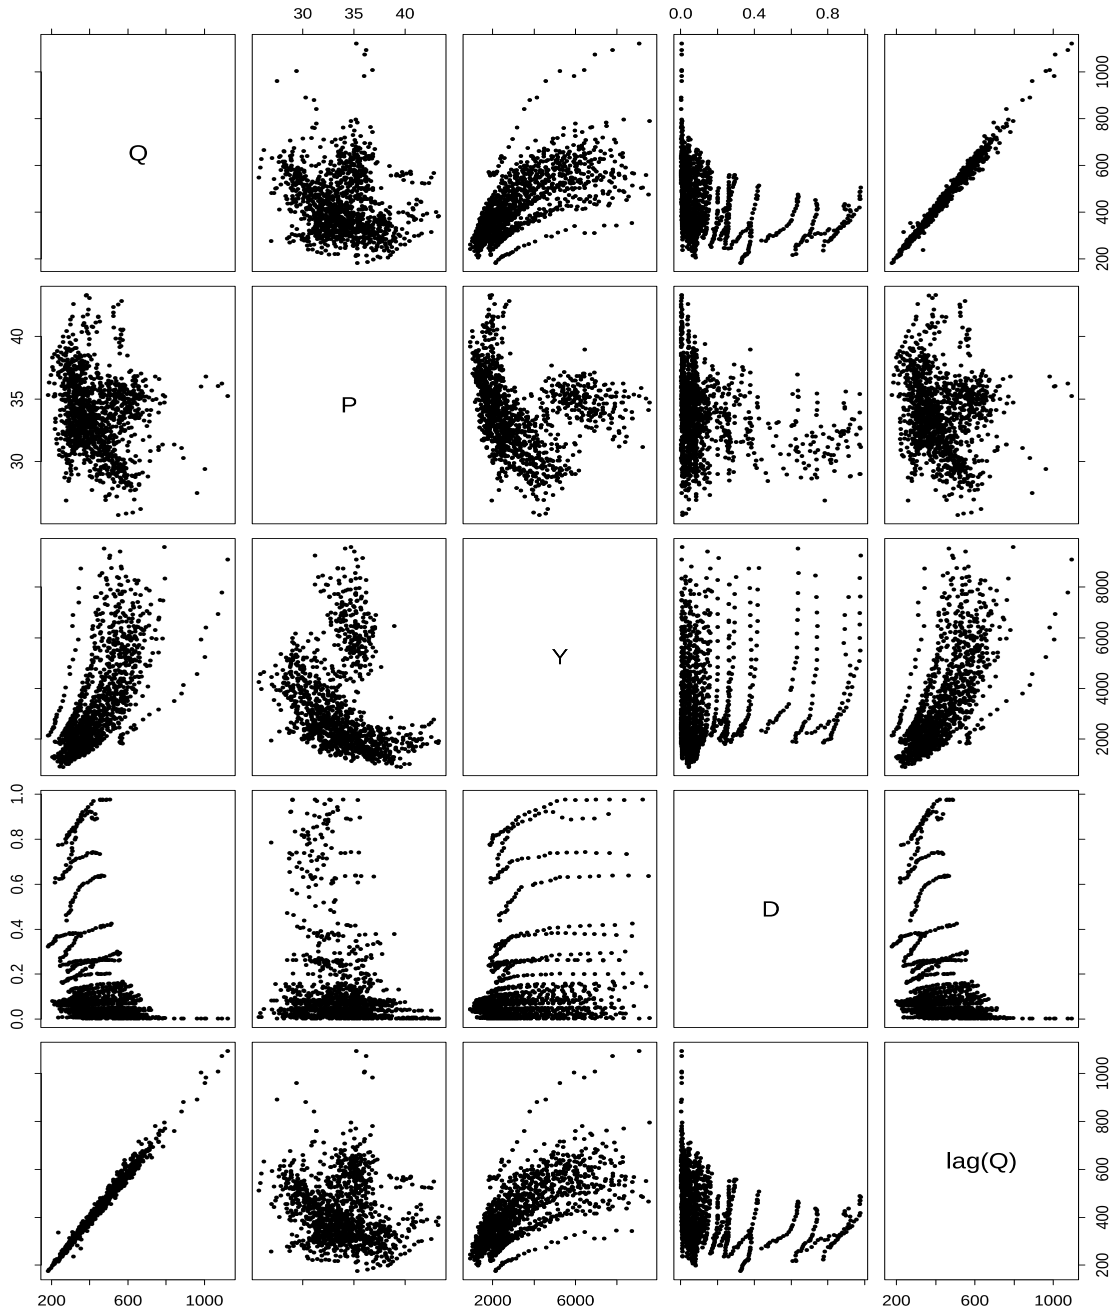

Figure 1: Scatter plots of raw data.

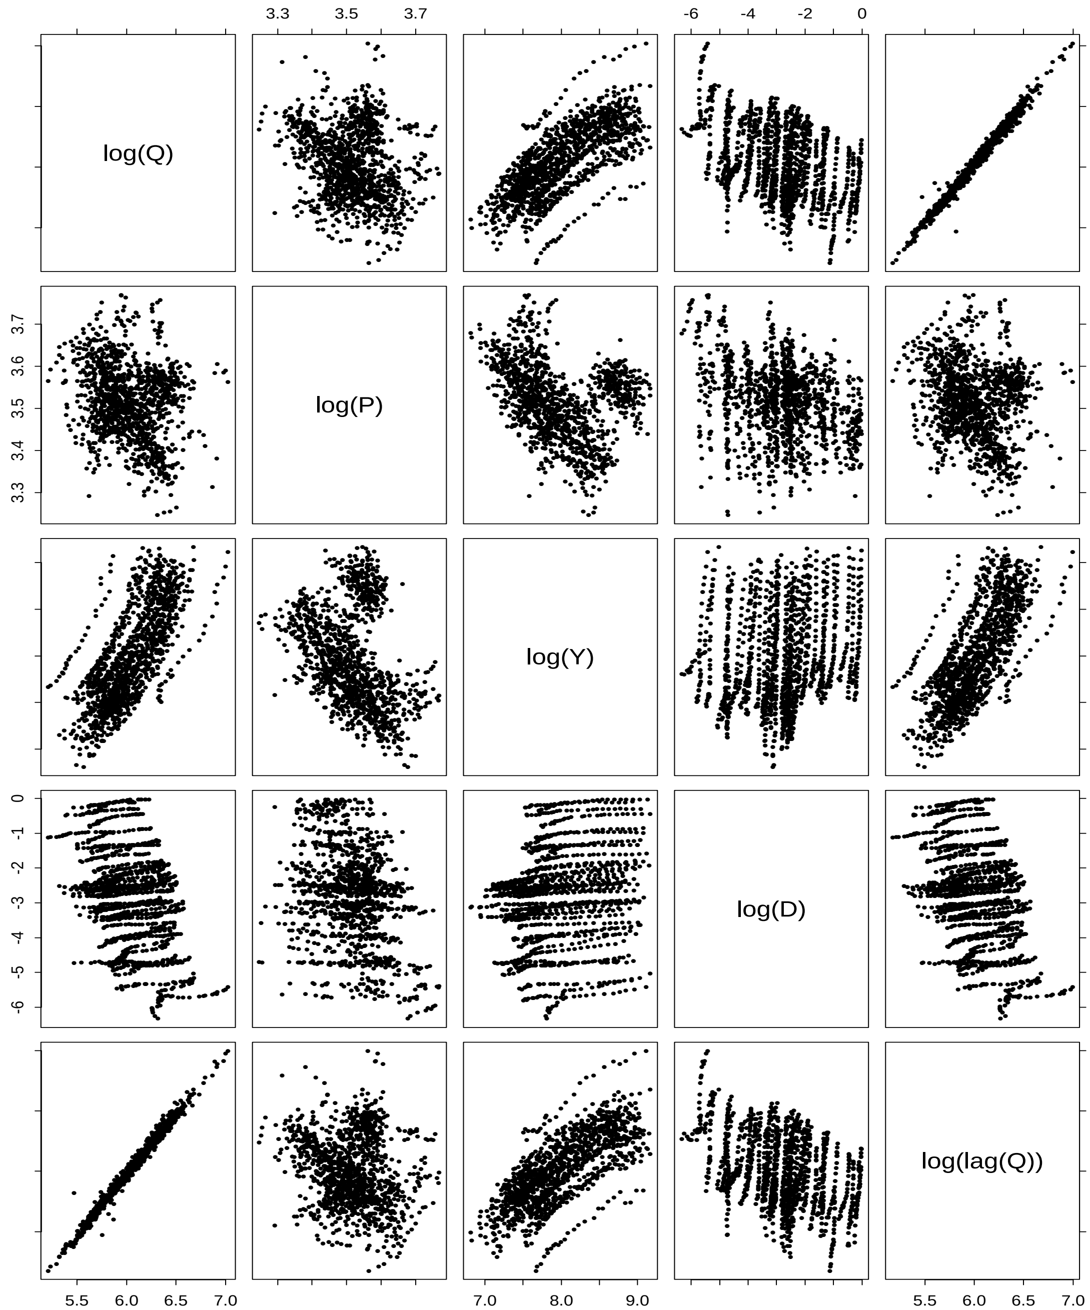

Figure 2: Scatter plots of logarithmic transformed raw data.

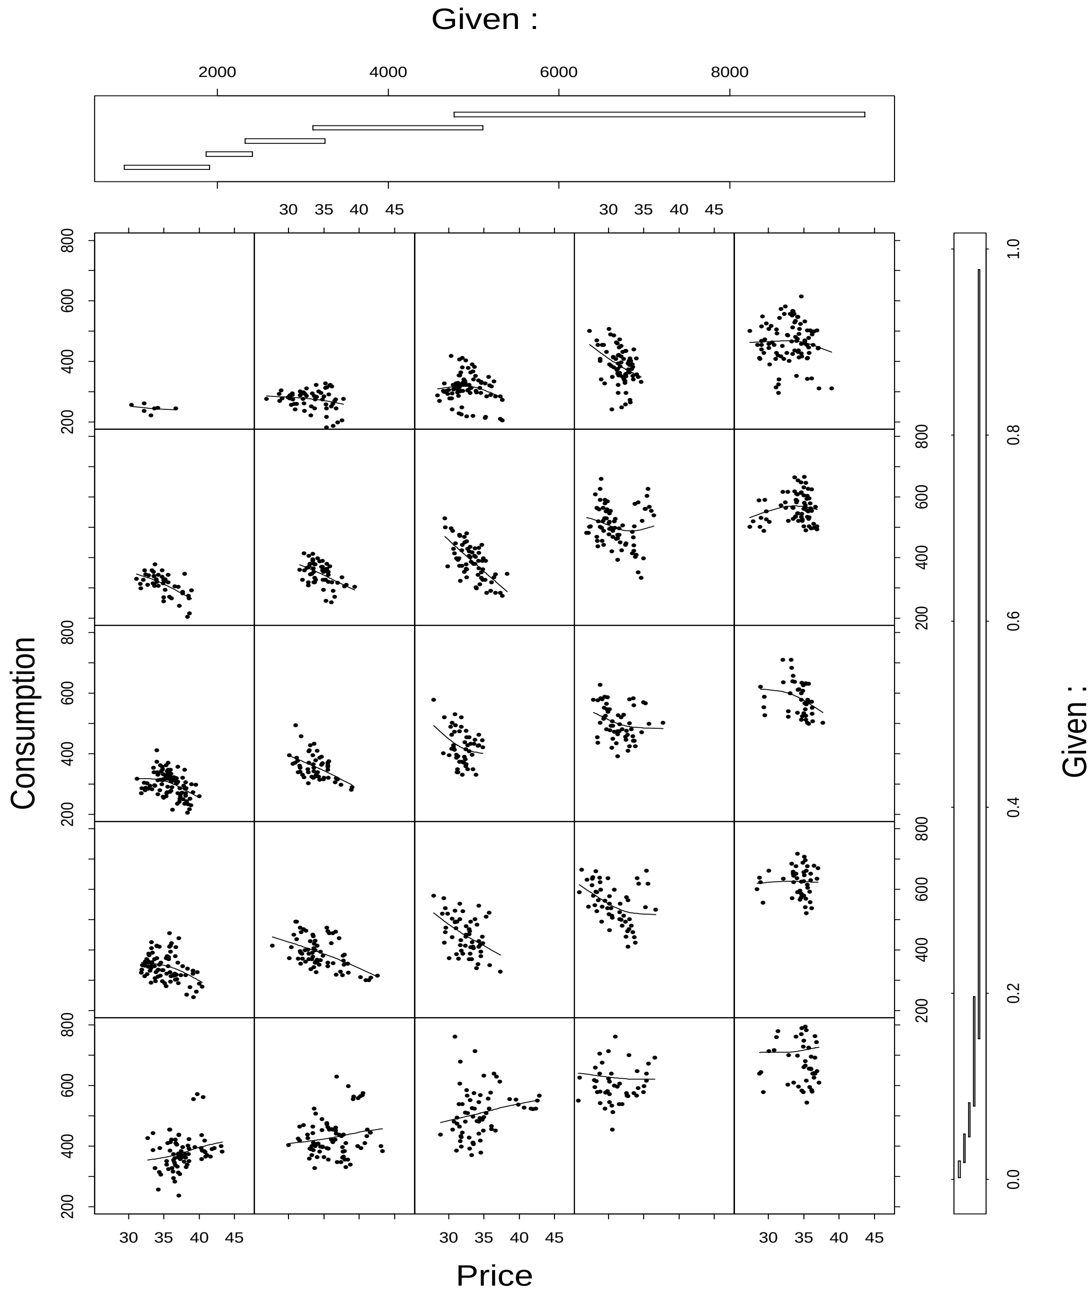

Figure 3: Conditional plots.

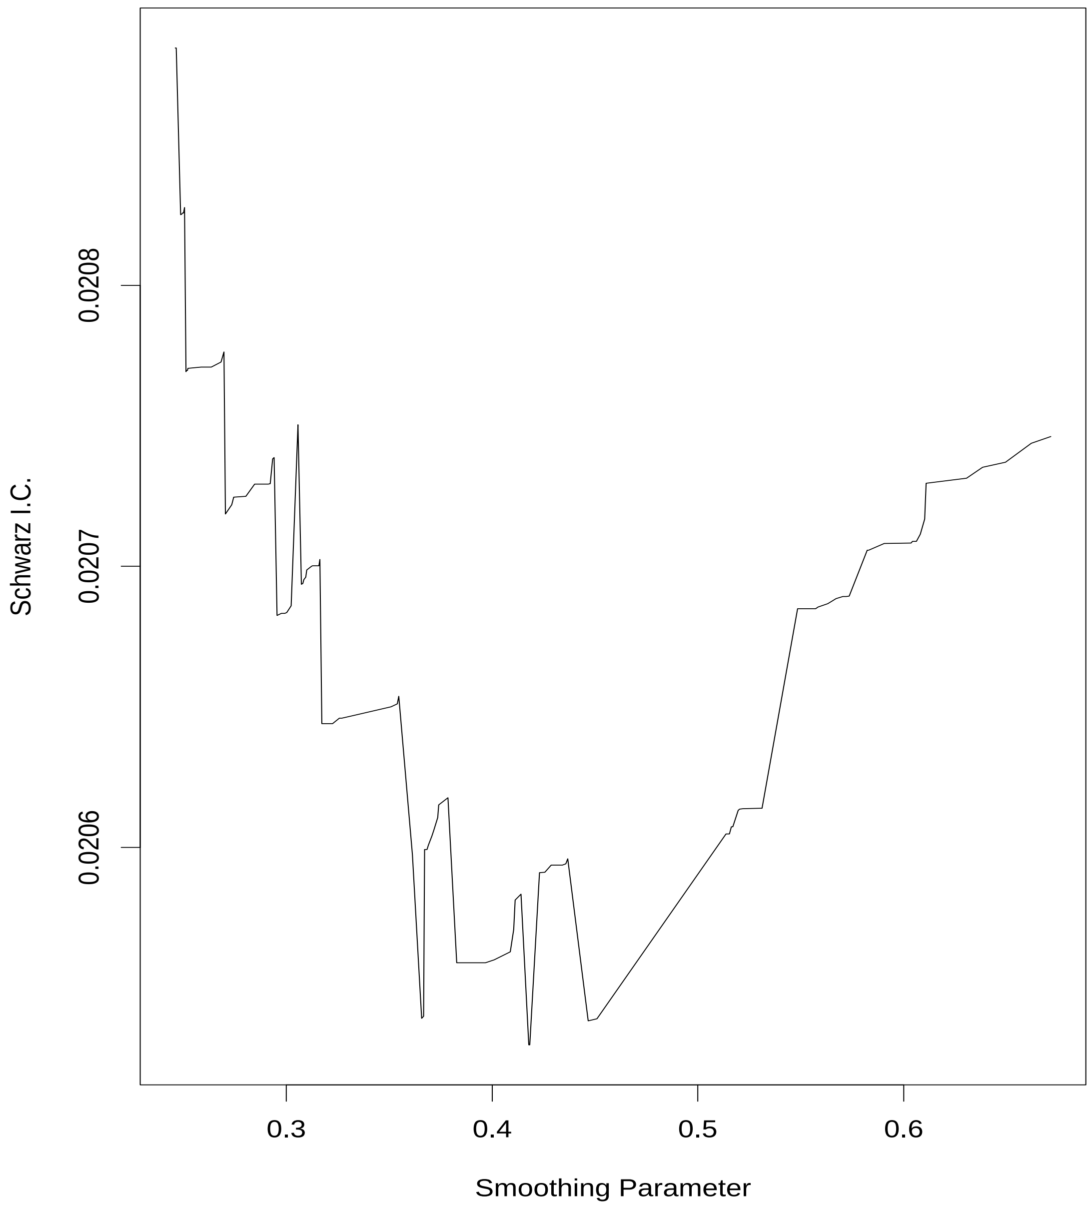

Figure 4: Schwarz information criterion.

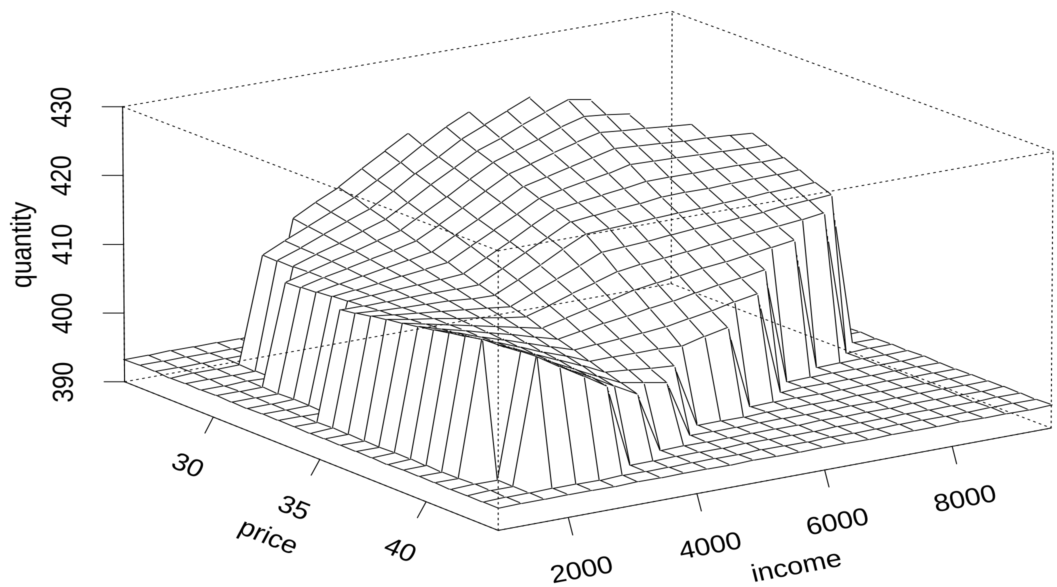

Figure 5: Smoothing B-splines demand surface.

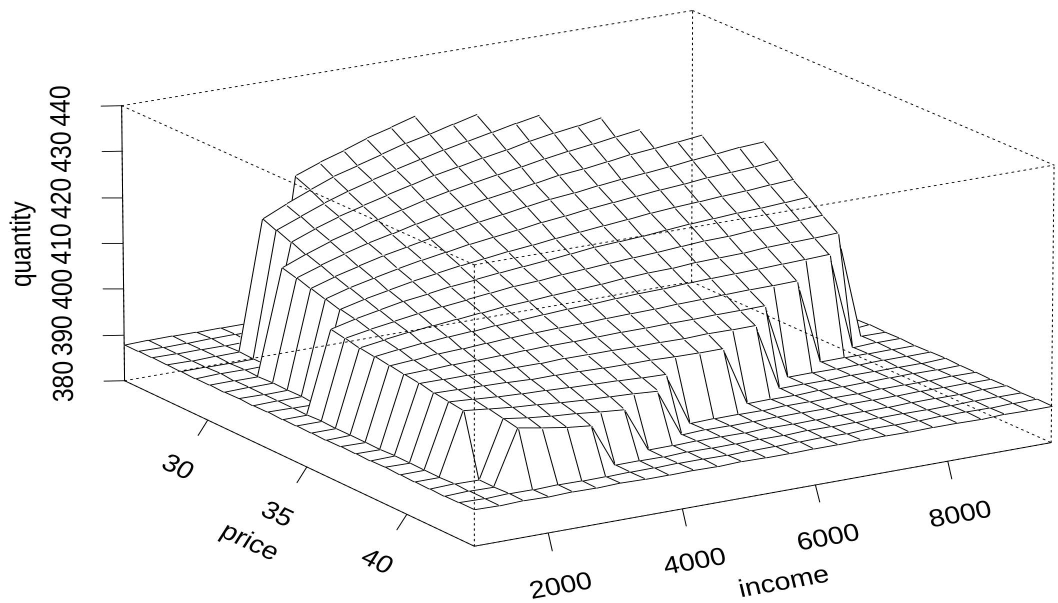

Figure 6: Translog Demand surface.

Figure 7: Smoothing B-splines confidence intervals for the demand curves: magnitude of price effect observed at increasing income levels.

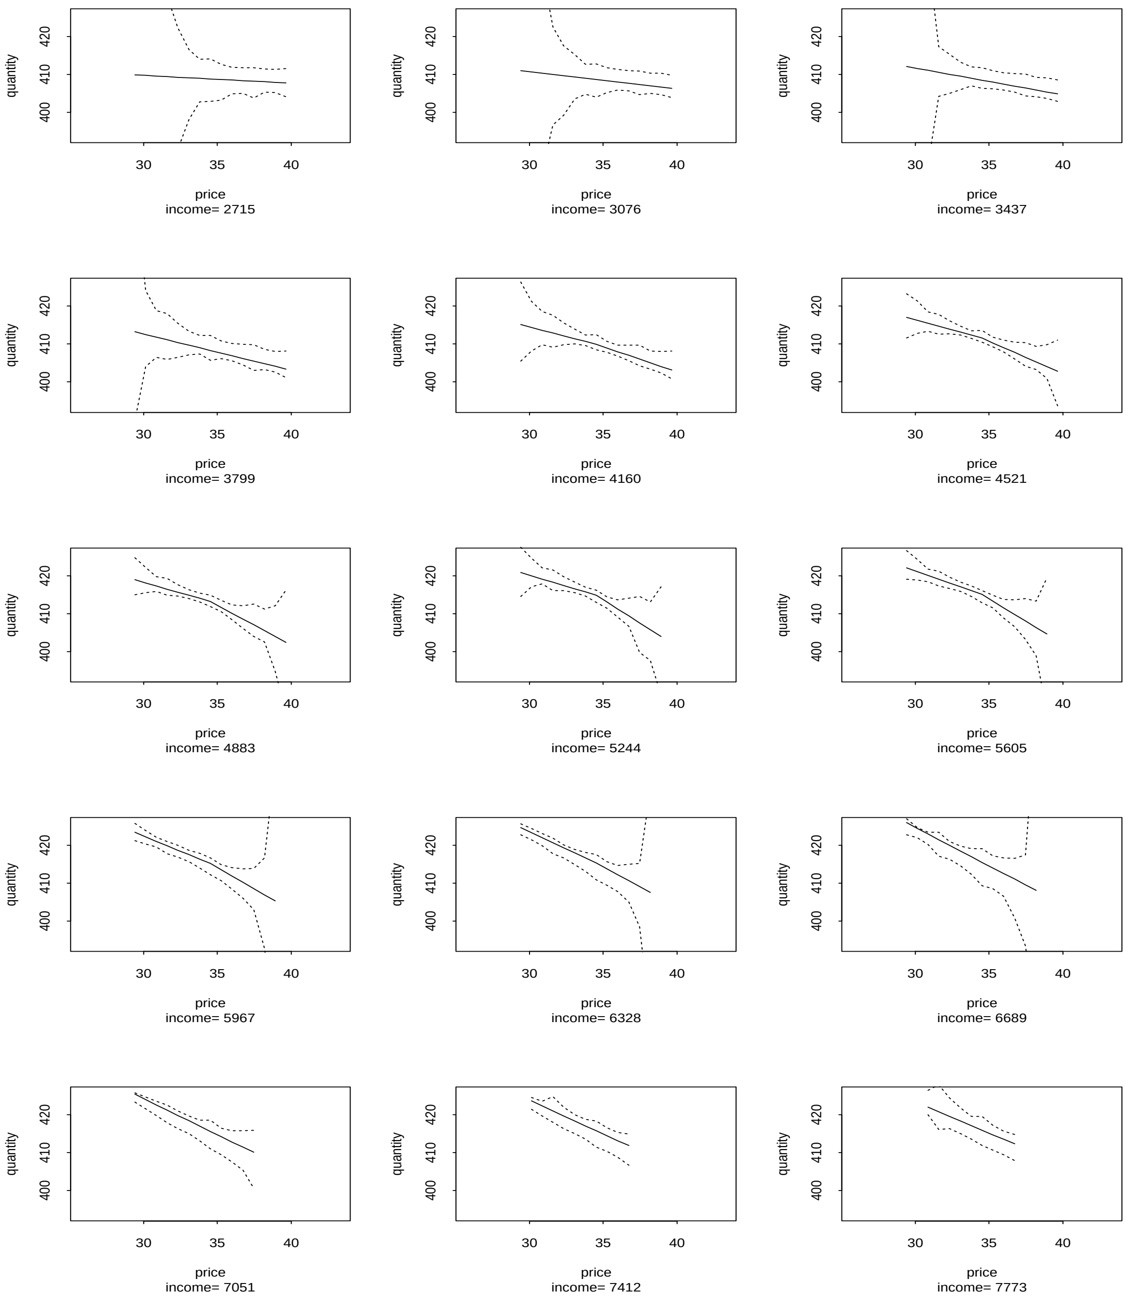

Figure 8: Translog Demand surface.

Figure 9: Confidence intervals for the translog demand curves: magnitude of price effect observed at increasing income levels.

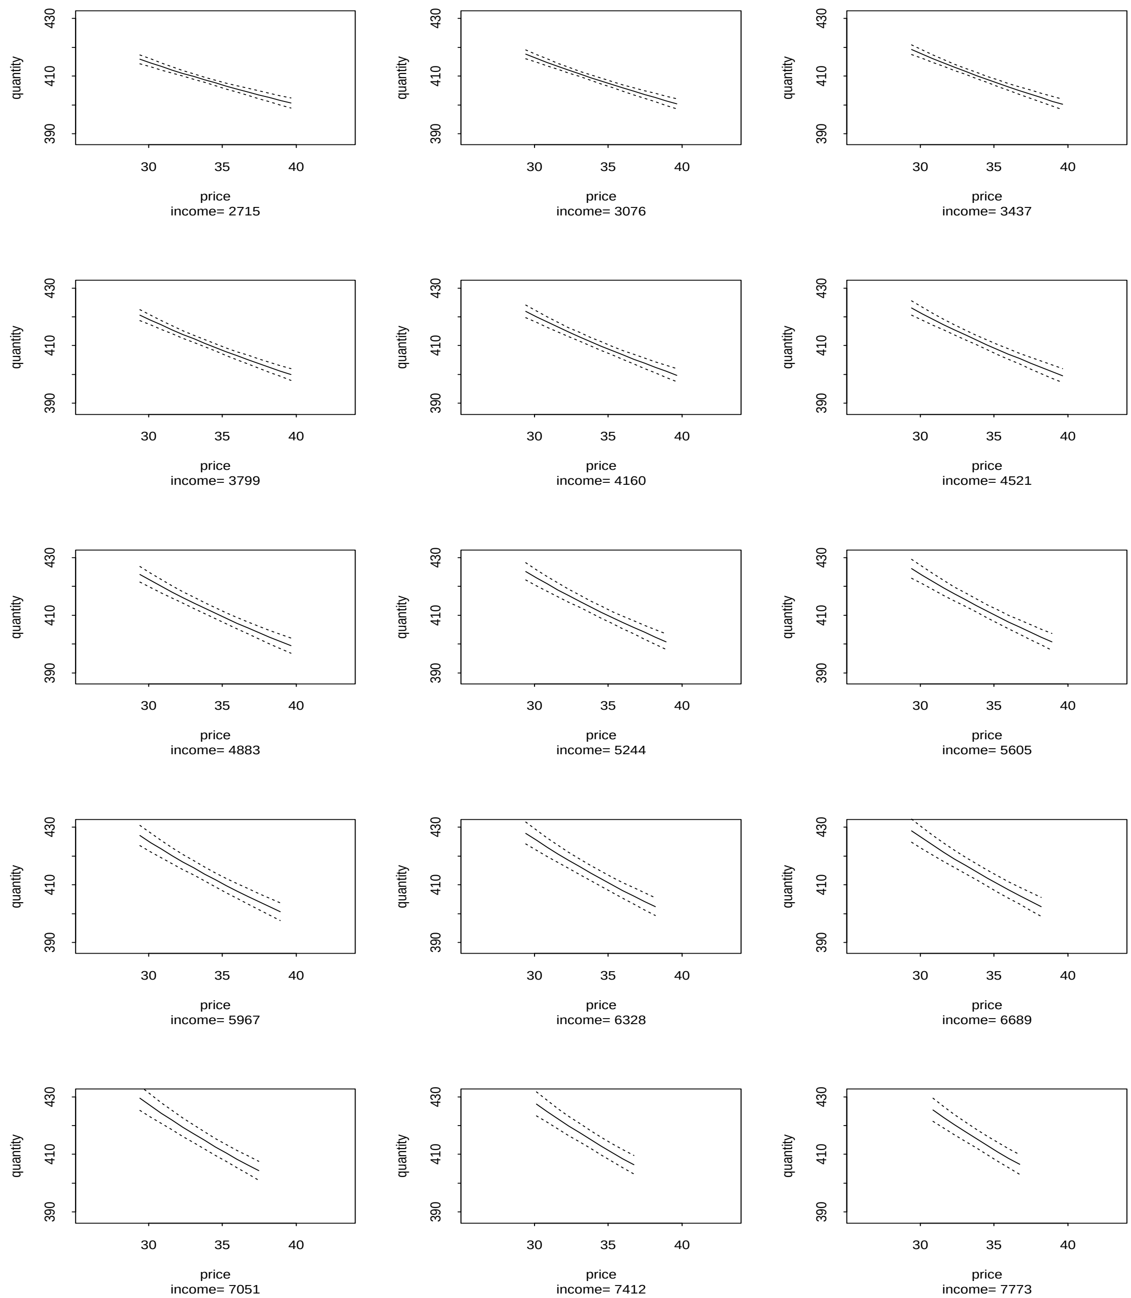

Figure 10: Translog Demand surface.

Figure 11: Smoothing B-splines short-run price elasticity.

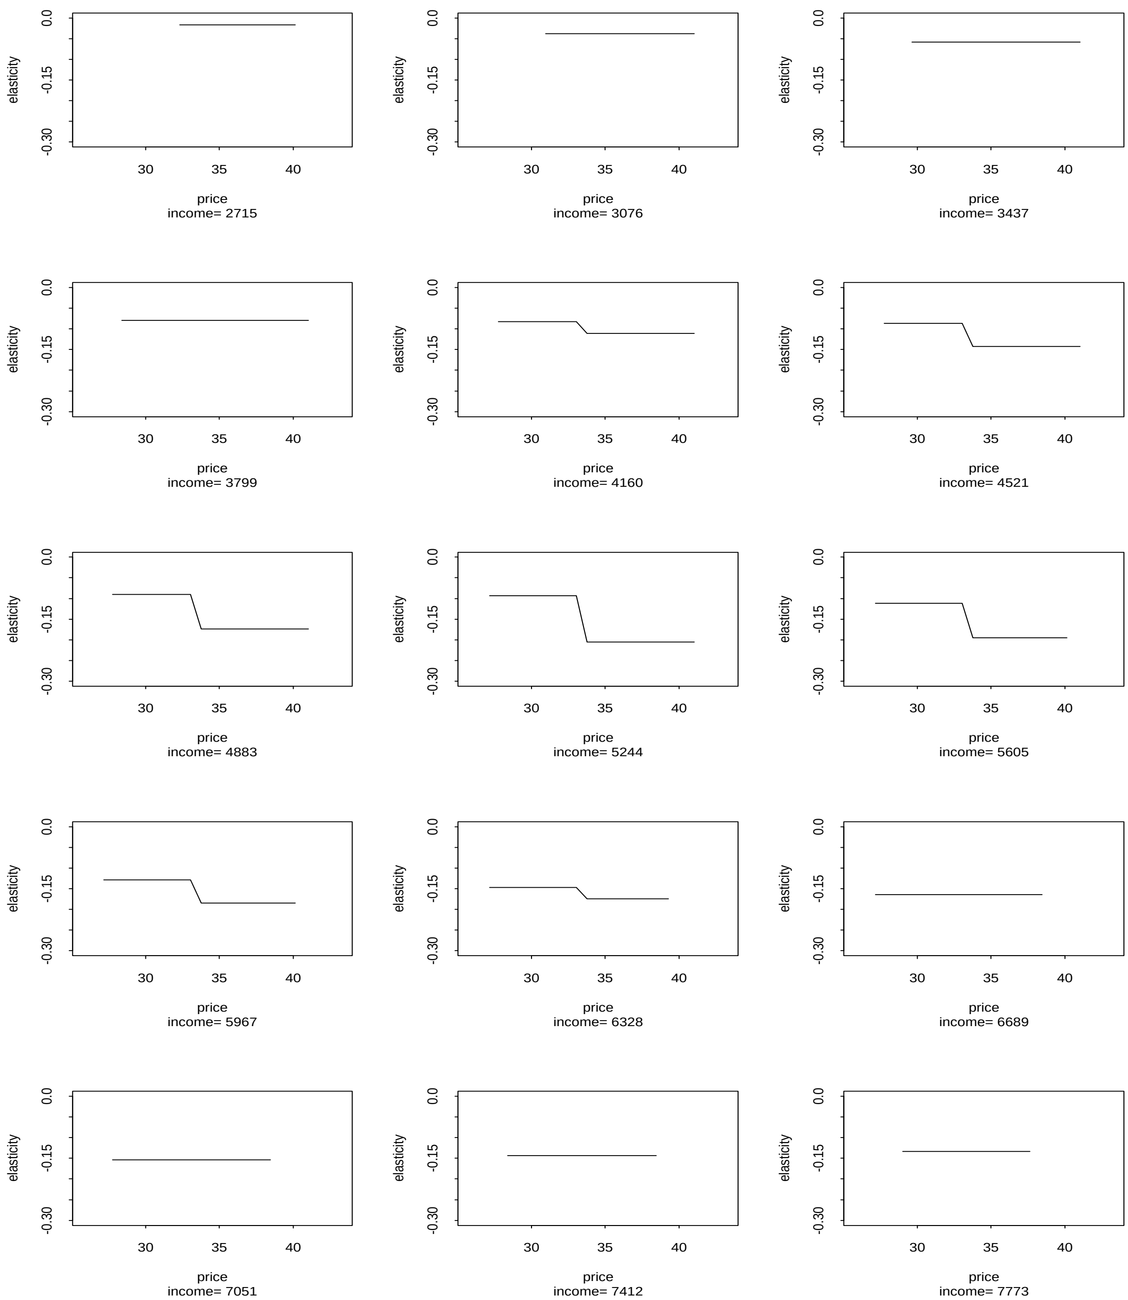

Figure 12: Translog Demand surface.

Figure 13: Short-run price elasticity for the translog model.

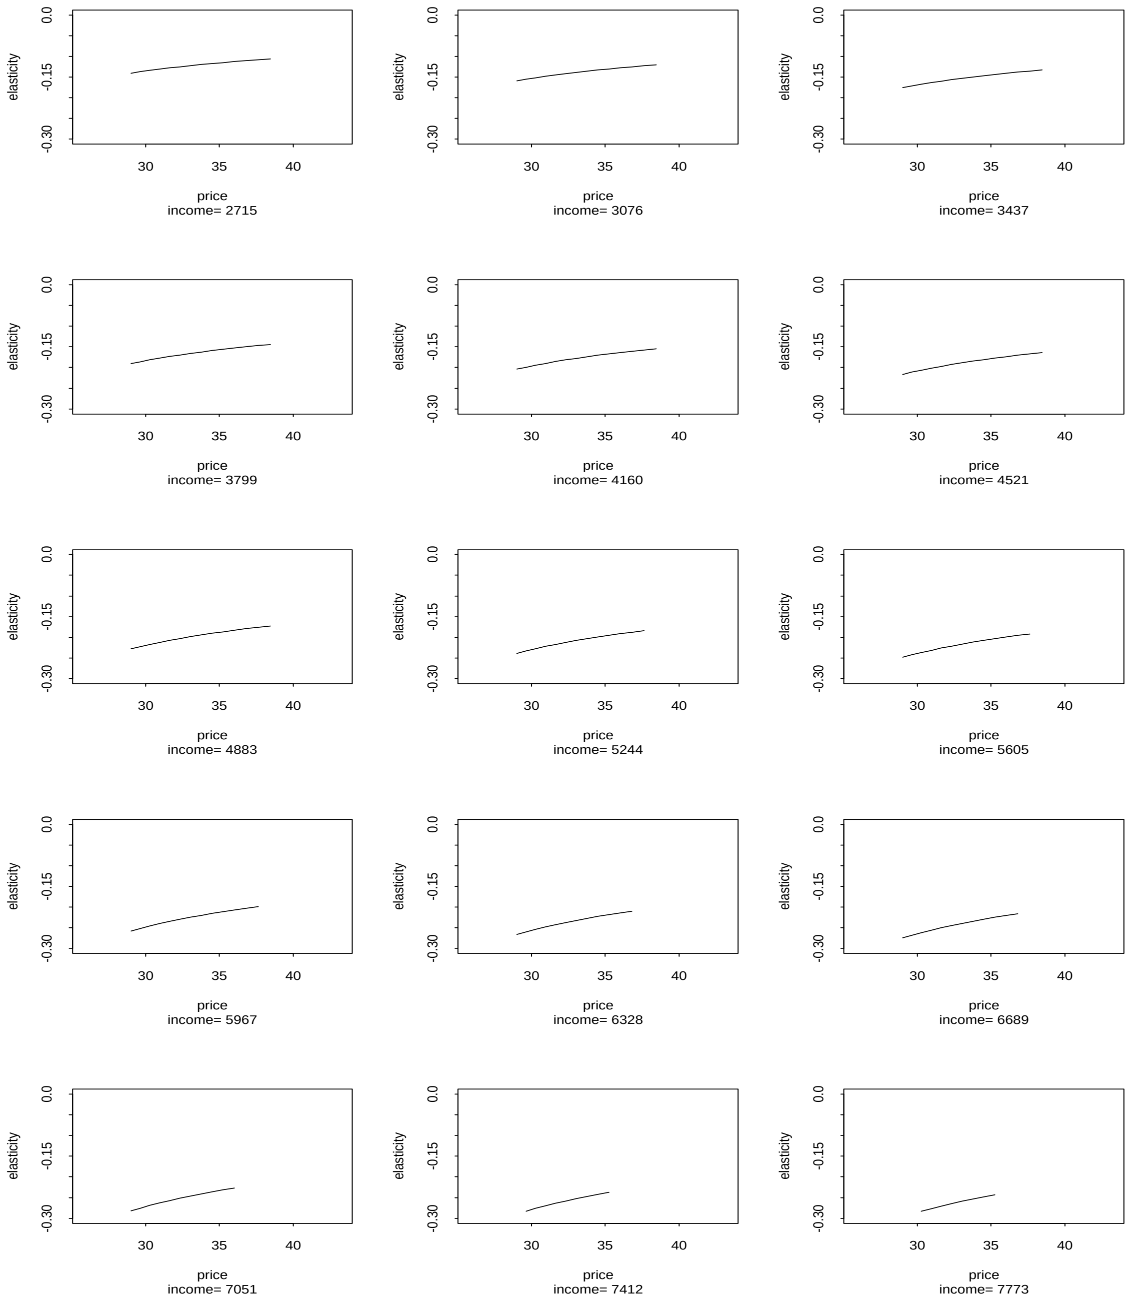

Figure 14: Translog Demand surface.

Figure 15: Smoothing B-splines confidence interval for the demand curves: magnitude of income effect observed at increasing price levels.

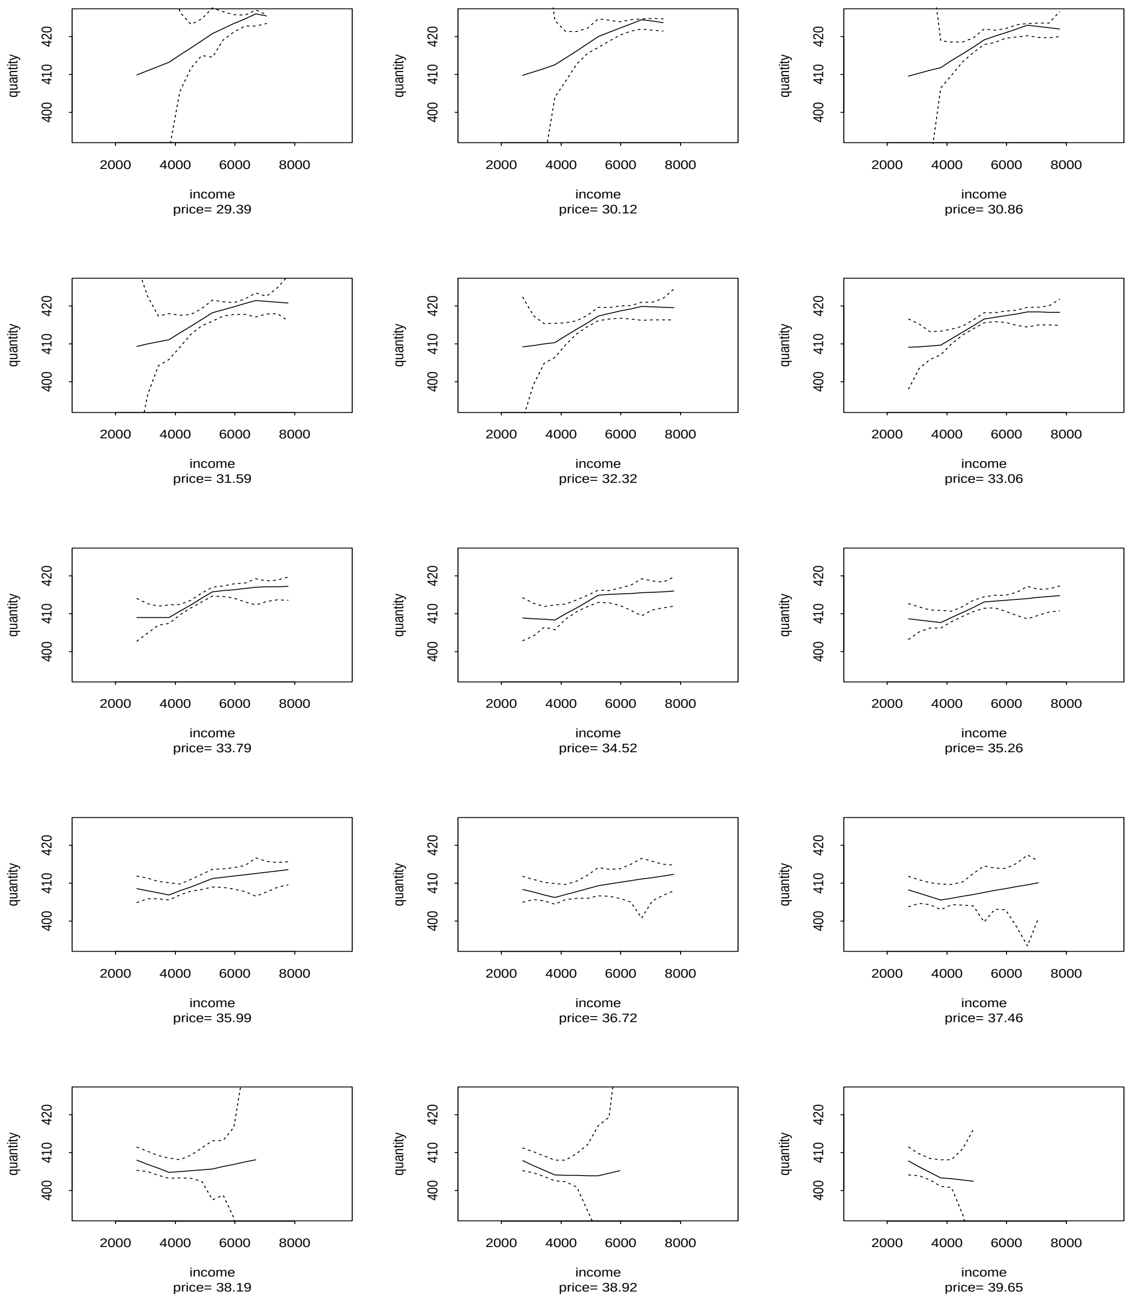

Figure 16: Translog Demand surface.

Figure 17: Confidence interval for the translog demand curves: magnitude of income effect observed at increasing price levels.

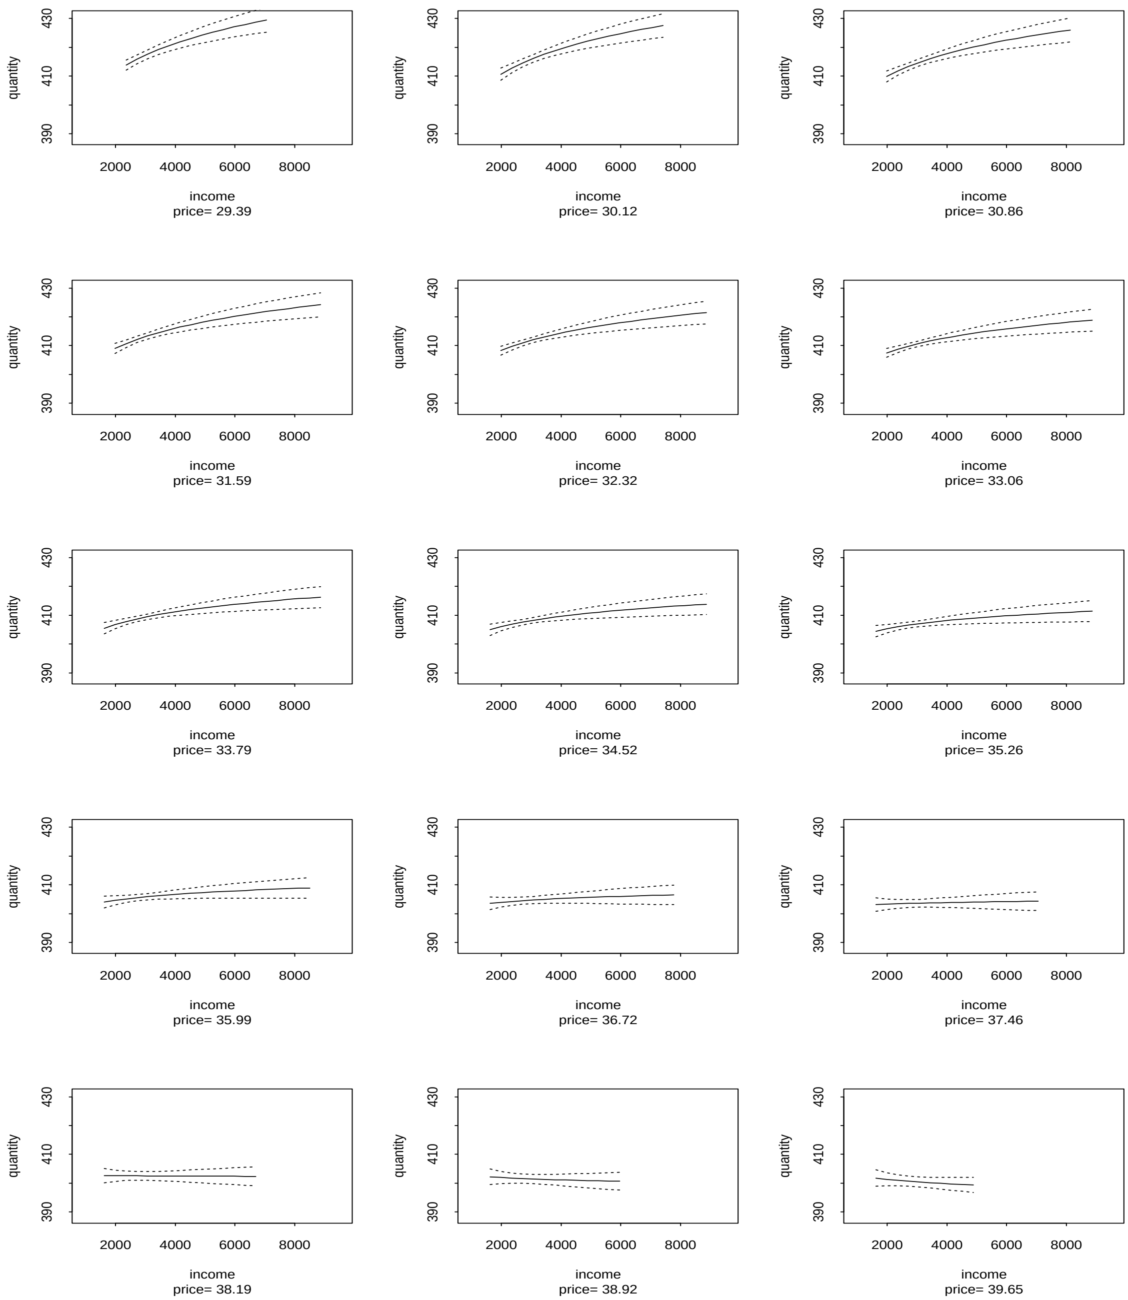

Figure 18: Translog Demand surface.

Figure 19: Confidence interval for the translog demand curves: magnitude of income effect observed at increasing price levels.

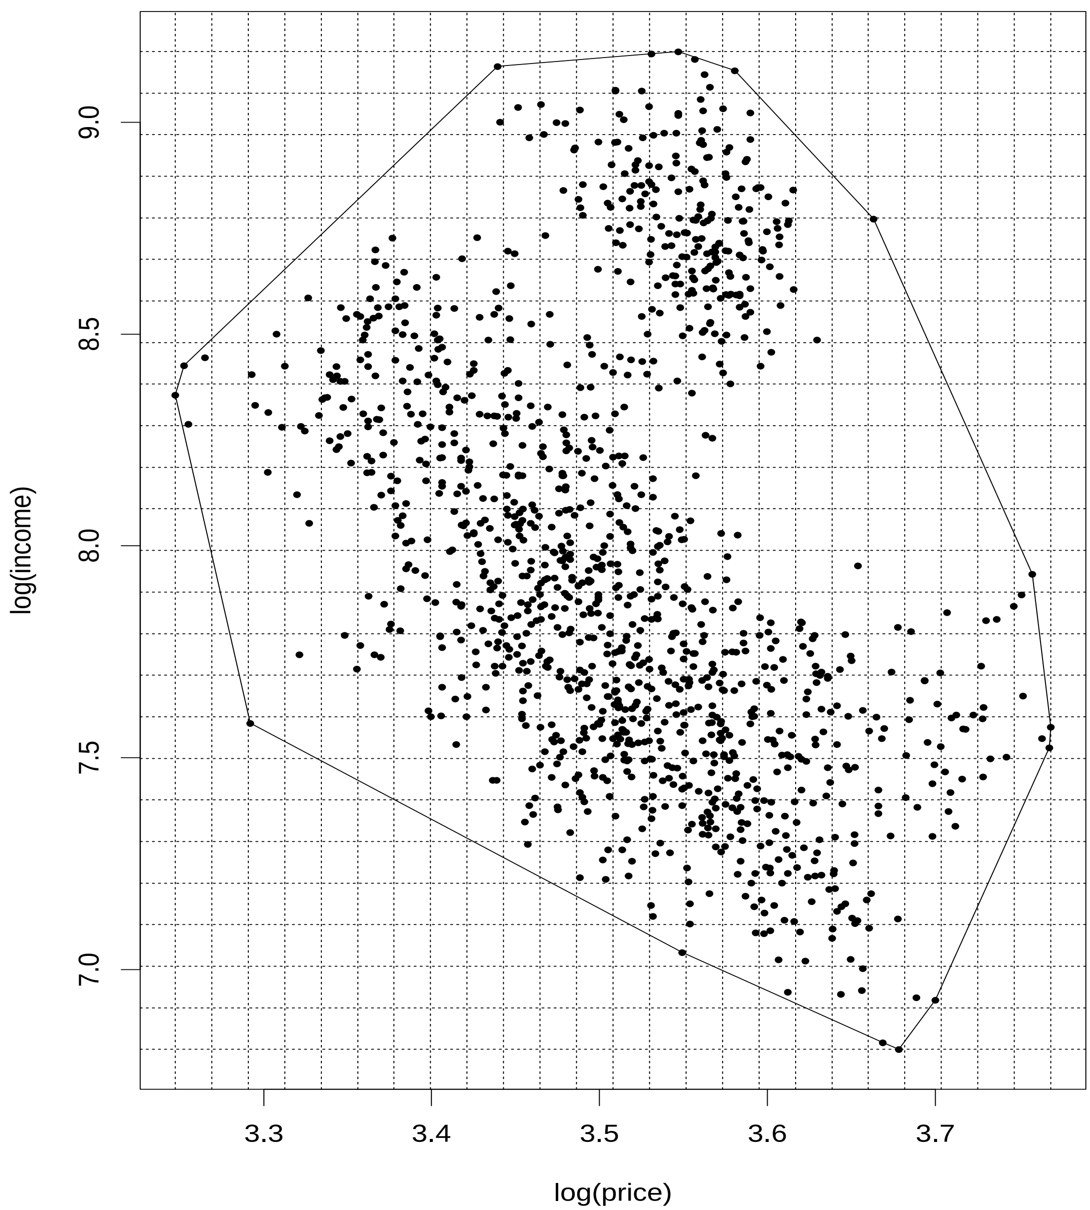

Figure 20: Convex-hull of data projected onto the log(price)-log(income) plane.

Figure 21: Smoothing B-splines short-run income elasticity.

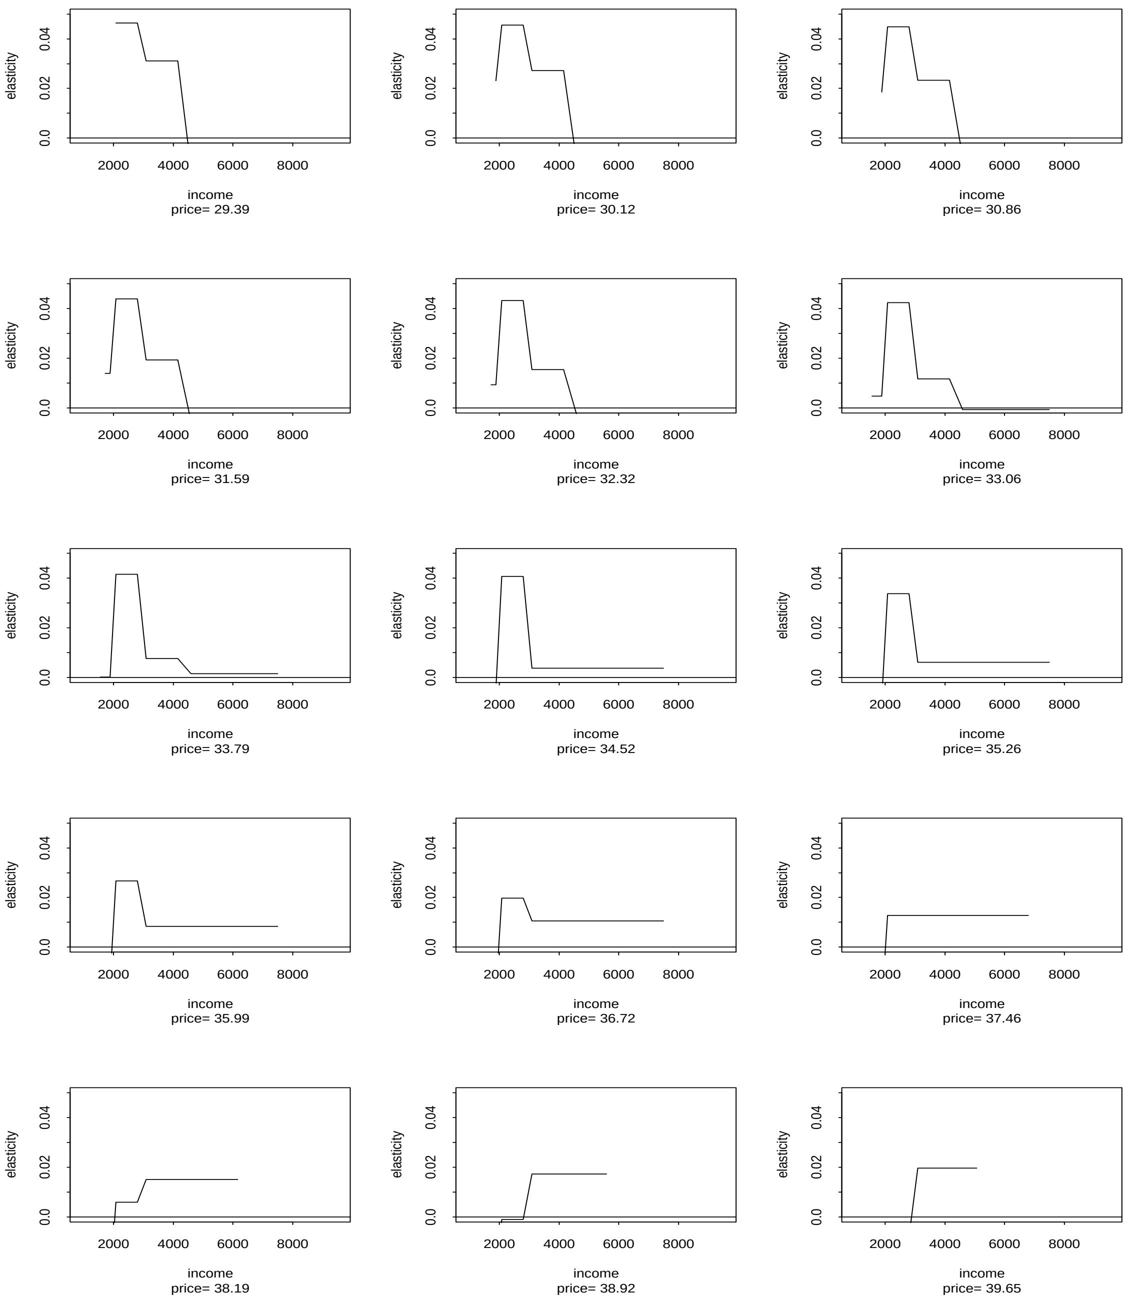

Figure 22: Convex-hull of data projected onto the log(price)-log(income) plane.

Figure 23: Short-run income elasticity for the translog model.

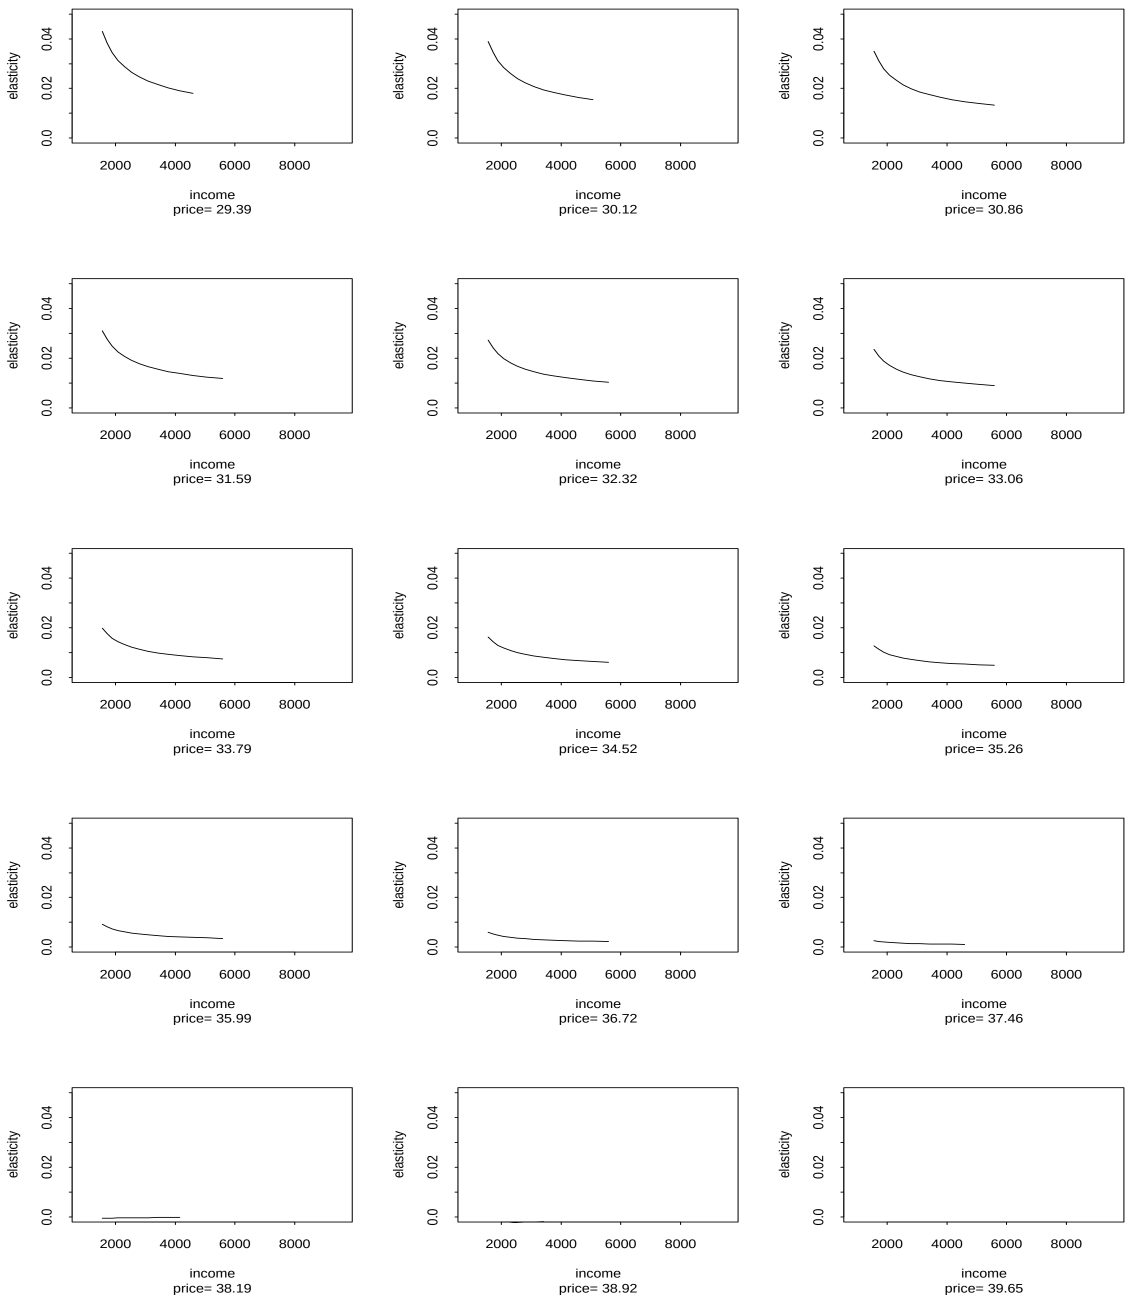

Figure 24: Convex-hull of data projected onto the log(price)-log(income) plane.