

Economic Development and the Structure of the Demand for Commercial Energy

Ruth A. Judson, Richard Schmalensee, and Thomas M Stoker*

To deepen understanding of the relation between economic development and energy demand, this study estimates the Engel curves that relate per-capita energy consumption in major economic sectors to per-capita GDP. Panel data covering up to 123 nations are employed, and measurement problems are treated both in dataset construction and in estimation. Time and country fixed effects are assumed, and flexible forms for income effects are employed. There are substantial differences among sectors in the structure of country, time, and income effects. In particular, the household sector's share of aggregate energy consumption tends to fall with income, the share of transportation tends to rise, and the share of industry follows an inverse-U pattern.

January 1998

* Judson: Board of Governors, U.S. Federal Reserve System, Washington, DC 20551. Schmalensee and Stoker: Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02139. The authors are indebted to the MIT Center for Energy and Environmental Policy Research for financial support and to Thomas Boden and his colleagues at the Carbon Dioxide Information Analysis Center at the Oak Ridge National Laboratory for data assistance. An earlier version of this paper was presented at the 1995 Meeting of the American Economic Association.

1. INTRODUCTION

This work is motivated by an interest in the relationship between economic development and the consumption of commercial energy. This relationship is clearly central to the development process itself. Moreover, carbon dioxide (CO₂) emissions produced by fossil fuel combustion are of independent policy interest because they may warm the earth's climate (IPCC (1996)). Forecasts of future CO₂ emissions depend primarily on explicit or implicit forecasts of the demands for commercial energy in its various forms and of the technologies that will be used to meet those demands.

In an earlier paper (Schmalensee, Stoker, and Judson (1998), hereafter SSJ), we used aggregate, national-level panel data for the 1950-1990 period to estimate reduced-form Engel curves for per-capita CO₂ emissions and commercial energy consumption. We allowed for country and time fixed effects and employed a flexible spline form for income effects. We encountered evidence of an "inverse-U" relationship between per-capita income and CO₂ emissions, along with weaker evidence of a similar relation involving energy consumption. In both cases the peak was estimated to occur below the sample maximum per-capita income. Inspection of the data indicated that the negative estimated income elasticities reflected declines in the carbon- and energy-intensity of OECD economies that began in the 1970s. Because our data covered 141 nations, we were able to obtain precise estimates of income elasticities for a wide range of development levels and thus to project global carbon dioxide emissions with reasonable confidence.

The objective of the present paper is to deepen understanding of the aggregate reduced form relation between economic development and energy demand by disaggregating our Engel curve analysis to the level of major economic sectors. Different sectors use different mixes of

energy sources, so that changes in the sectoral composition of GDP generally have implications both for the structure of energy demand and for emissions of CO₂ and various pollutants. In addition, future advances in technology are likely to have differential impacts on energy consumption in different sectors. Thus, for instance, transportation mainly uses petroleum products, and its demand for them will be little affected by advances in heating and cooling of buildings or (absent dramatic improvements in battery technology) in electricity generation. Finally, national differences both in development at the sectoral level and in the sectors' shares of economic activity may have important implications for current and future energy consumption.

Since the bulk of the world's population is outside the OECD region, and since it seems clear that future increases in energy consumption and carbon dioxide emissions are likely to come primarily from non-OECD nations,¹ we felt it was necessary to analyze data for both OECD and non-OECD nations. As in our earlier work, however, including significant coverage of non-OECD countries immediately ruled out estimation of price effects, because comprehensive data on domestic energy prices in non-OECD nations do not exist.² The good news is that the United Nations compiles rich and comprehensive data on sector-level energy consumption that (but for the lack of price information) seem on the surface ideal for our purpose. The bad news is that these data appear to be infected with unusual severity by measurement error. As Section 2 discusses, however, we believe we have excluded the most

¹ This point is discussed and related references are provided in SSJ.

² We were reluctant to restrict our sample to include only the few non-OECD nations for which price data were available, since their being unusual in this regard may signal that they are unrepresentative in other relevant respects as well. Moreover, one can argue that it would in any case be inappropriate to treat domestic prices as exogenous in a study concerned primarily with long-run relationships. World market prices determined endogenously, and domestic energy prices are determined to an important extent by tax and subsidy policies, which are shaped by relations between economic development and political behavior.

suspect observations and that the estimates reported here are based on a data set with a great deal of information on sectoral energy usage.

Section 2 also describes our estimation methods. With data covering a wide range of per-capita incomes, it seems unjustified to assume that the income elasticity of sectoral demand for energy is constant. We employ here the spline functional form used in SSJ, where we found the income elasticity of aggregate energy demand to vary significantly and substantially with income. In addition, we make use of the relatively good data on national “apparent consumption” (defined below) to mitigate the effects of remaining measurement problems.

Section 3 presents our estimation results. While country effects are highly correlated across sectors and with per-capita GDP, there are substantial differences among major sectors in patterns of time and income effects. Over our 1970-1991 sample period, time effects account for considerably more rapid growth in the household and agricultural sectors than elsewhere. (No sector’s time effects are highly correlated, with either sign, with world oil prices.) This plausibly reflects the international diffusion of such energy-using innovations as televisions, tractors, and air conditioners. As per-capita income rises, all else equal, our estimates imply a decline in the household sector’s share of energy consumption and an increase in the share of transportation—particularly at the highest income levels. The share of the industry and construction sector is estimated to follow an “inverse-U” pattern with income growth. Most of our estimated income elasticities are below unity.

There is relatively little prior work to which this study can be directly compared. Most studies of energy demand at the sector level have focused on OECD nations; see Pindyck (1979) for an early and still instructive example. In part because of data limitations, the relatively few sector-level studies of non-OECD energy demand tend to concentrate on single

sectors and/or single nations; see the surveys by Bhatia (1987) and Dahl (1992). Most previous work has not employed samples that made it possible to distinguish time, country, and income effects, and restrictive functional forms have generally been employed to model the effects of changes in income.

2. DATA AND METHODS

The analysis in SSJ was based on United Nations (UN) data on national “apparent consumption” of fossil fuels: (imports + production) minus (exports + stock increases). Imports and exports of fuel are generally tracked with some care, and significant domestic energy production usually attracts significant government attention. Thus imports, exports, and production are generally measured fairly accurately, and these are generally the most important determinants of total apparent consumption, particularly over substantial spans of time. Accordingly, the UN apparent consumption statistics and estimates of carbon dioxide emissions based on them are generally considered reliable.³ In contrast, sectoral consumption figures in the UN data are based on direct estimates, either by national governments or by the United Nations staff. Because individual nations generally have less interest in sectoral consumption levels than in imports and exports, and many countries accordingly lack mechanisms to measure energy consumption accurately on a disaggregated basis, sectoral consumption estimates are inherently less reliable than aggregate estimates.⁴ This section

³ See Marland et al (1989) and Alcamo et al (1994).

⁴ Sectoral consumption data are also produced by the International Energy Agency (IEA). It is our understanding that the IEA and UN data are based on the same information from national governments -- collected for OECD nations by the IEA and for other nations by the UN. Nonetheless, the published figures differ, sometimes substantially. We chose to work with the UN data because they link to the generally accepted carbon dioxide emissions statistics and because we suspected that having apparent consumption data from the same source would be useful in dealing with problems of measurement error.

outlines how we dealt with this problem in constructing our data set and concludes with a discussion of our estimation methods.

Definitions and Sources

The UN sectoral consumption data came in the form of 105,718 observations covering the 1950-1992 period.⁵ Each observation contained year, fuel, country, and sector codes, along with a quantity measure that we converted to Btus using conversion factors furnished with the data. We eliminated about 10 percent of these observations as follows. We dropped all observations before 1970 because there were only 457 of them, and we dropped the 3133 observations for 1992 because (as noted below), we had no income data for that year. In principle the data contain 54 fuel codes, but there were no observations on 14.⁶ An additional seven fuel codes that accounted for 5391 observations were dropped. Three of these were traditional fuels, and all had spotty coverage and raised serious data quality concerns.⁷ These deletions narrowed our focus to commercial energy consumption, though we drop “commercial” in what follows for brevity. Finally, two countries were dropped because they reported no consumption of any of the remaining 33 fuel types.

Our income and population data were all taken from the Penn World Table, Mark 5.6, which covered the period 1950-1991.⁸ We used the RGDPH series for GDP. This series is based on a chain index of prices in each country and employs

⁵ For descriptions of these data and definitions of sectors, see United Nations (1994a), (1994b), and (1995). We are indebted to the Carbon Dioxide Information Center at the Oak Ridge National Laboratory for providing the UN data to us in a convenient form.

⁶ The corresponding codes are MW, OS, NC, PU, TH, UR, GL, MP, WF, EG., EH, EN, EW, ET.

⁷ After discussions with the Carbon Dioxide Information Center at the Oak Ridge National Laboratory, we dropped all observations for the following: animal wastes, bagasse, fuelwood, vegetal wastes, alcohol, and blast furnace gas. In addition, we dropped the 19 observations involving plant condensate.

⁸ See Summers and Heston (1991) for a description of these data.

estimates of purchasing-power-parity exchange rates in 1985 to convert all GDP figures to 1985 U.S. dollars. Of the 187 countries remaining in the United Nations sample, 37 were dropped because we had no GDP data for them, and an additional 5 were dropped because GDP data were available only for 1985. Three additional country codes were dropped as a consequence of boundary changes. At this stage, the United Nations data consisted of 91,615 sector-level observations on 33 fuel types in 142 countries over 22 years. Adding across fuel types gave estimates of sectoral energy consumption.

We computed energy consumption in this fashion for the five major final demand sectors in the UN data: *Industry and Construction*, *Transportation*, *Households and Other*, *Energy Sector*, and *Non-Energy Uses*. (The latter includes use of petroleum products as inputs in chemical processes and as road paving.) In addition, we computed consumption for two potentially interesting sub-sectors of *Households and Other*: *Households* and *Agriculture*.⁹ In the absence of gaps in the UN data, we would have had 3124 (=142 x 22) observations on each sector or sub-sector. In fact, there were significant gaps, and we had from 2740 (*Households and Other*) to 1515 (*Energy Sector*) observations available.

We then carried out an observation-by-observation inspection of these data. We uncovered many implausible values that seemed to indicate measurement problems. In particular, absurdly large year-to-year changes in estimated sectoral consumption were not uncommon, often accompanied by changes in the number of fuels measured. In other cases, estimated consumption would be constant to four digits for several

years. The *coverage ratio*, the ratio of total estimated sectoral consumption to total apparent consumption, was often quite small, particularly when only a few fuels were reported having been measured, and this ratio sometimes varied substantially over time.

Given the size of our samples, the danger of including a few observations with huge measurement errors seemed to us to outweigh the problem of dropping a few valid observations. We accordingly we elected to make a systematic effort to exclude suspicious observations. We proceeded in two stages.

At the first stage, two of us went through the data independently and marked observations that seemed suspicious. Our judgments were informed by indicators of three sorts of year-to-year changes:¹⁰ (1) sectoral consumption was less than 50 percent or more than 150 percent of the previous observation, (2) the ratio of sectoral consumption to total apparent consumption changed by 10 or more percentage points or was less than 50 percent or more than 200 percent of the previous observation, and (3) the number of fuels covered by the sectoral estimate changed.¹¹ When changes of types (1) or (2) separated blocks of data, the block with the more comprehensive fuel coverage and/or the most recent data was generally retained. Data for 1990 and 1991 often had less comprehensive fuel coverage than data for earlier years, and, where changes of types (1) or (2) were present, they were often dropped for that reason. In

⁹ The omitted sub-sectors are *Public* and *Other Consumers*.

¹⁰ We spent some time trying without success to devise an algorithm that would generally agree with our judgments as to which observations should be dropped.

¹¹ It has been suggested to us that we could avoid the problem of time-varying coverage of energy sources by limiting our attention to a single important source: electricity. We are following that suggestion in ongoing work.

addition, blocks of data in which one or two four-digit consumption estimates were repeated for several years were dropped, particularly when the coverage ratio was low.

At the second stage in this process, we dropped almost all observations that either of us had marked as suspicious or for which the coverage ratio was less than 0.20, along with a few observations that were the only remaining ones for the corresponding countries. Between 27 percent (*Households*) and 34 percent (*Non-Energy Uses*) of the available observations were dropped in this fashion. Finally, from 1.4 percent (*Non-Energy Uses*) to 3.6 percent (*Industry and Construction*) of the remaining observations were dropped because of lack of GDP data.

Data Overview

Table 1 provides information on our final data set. While a substantial number of non-OECD nations were dropped entirely, all samples include at least 69 countries and over half of world population in 1980, and at least 61 percent of all sectoral samples are composed of data from non-OECD nations. Table 1 indicates that there were also problems with the UN data on OECD member states. Even though no observations from any of the 24 sample-period OECD nations were dropped because of lack of GDP data, at least 80 OECD observations in each sector were dropped because of implausible year-to-year movements. (Otherwise there would have been $24 \times 22 = 528$ OECD observations on all sectors.)

As the per-capita and per-GDP means in Table 1 indicate, the first three sectors listed in the Table are clearly the most important. Moreover, energy consumption by the *Energy Sector*, for which our initial and final samples are smallest, is likely to reflect differences in resource endowments and international market factors that we do not

measure at least as much as it reflects the development of the domestic economy. Accordingly, in the interest of brevity, the discussion in what follows concentrates primarily on results for *Industry and Construction*, *Transportation*, and *Households and Others*. The samples for these three main sectors cover at least 119 nations and 69 percent of world population in 1980, and they have at least 76 percent of their observations from outside the OECD.

The last few columns in Table 1 are intended to give a rough, qualitative impression of the pattern of sectoral energy demand internationally. Putting aside income and time effects, if the sectoral income elasticity of demand for energy tends to be closer to unity than to zero, one would expect energy consumption per dollar of GDP to have a smaller coefficient of variation than consumption per capita. Table 1 reveals that this pattern generally holds, most clearly for *Transportation*.

Methodology

The basic reduced form Engel curve model we estimate in this study is an extension of that used in Schmalensee, Stoker, and Judson (1998):

$$C_{sit} = \alpha_{si} + \beta_{st} + f_s(Y_{it}) + \rho_s R_{it} + \varepsilon_{sit}, \quad (1)$$

where C_{sit} is the log of per-capita consumption in sector s in country i in year t , α_{si} is the country effect for country i in sector s , β_{st} is the time effect for year t in sector s , Y_{it} is the log of per-capita real GDP in country i in year t , ρ_s is a constant specific to sector s , R_{it} is the log of the coverage ratio for country i in year t , defined above, and ε_{sit} is a disturbance term with the usual properties. As noted above, real domestic prices do not appear in this equation because there are no data from which they could be computed on a systematic, comprehensive basis.

The R_{it} term in equation (1), which does not appear in our earlier work, is designed to reduce the impact of certain kinds of measurement error on the estimated income response function, f_s . In the absence of measurement error, the coverage ratio would be approximately unity, and R_{it} would thus be approximately zero.¹² This variable accordingly has no effect for “clean” observations. Otherwise, R_{it} captures systematic shortfalls in C_{it} due to incomplete measurement, so that the ρ_s are expected to be positive, and inclusion of the coverage ratio can be expected to improve income elasticity estimates.¹³ Increases over time in the fuel coverage of national energy statistics, a not uncommon pattern, will generally cause the coverage ratio to rise. Since per-capita income also tends to rise over time, leaving R_{it} out of the estimating equation would likely tend to produce an overestimate of the effect of income on demand.

As in our earlier related work, we employed a flexible spline (piecewise linear) form for the income response functions, f_s . We began with 24-knot splines, using the same locations of the knots (kinks) for all sectors to facilitate comparability. These locations were chosen by dividing the range of per-capita GDP data so as to yield the same number of potential observations in each inter-knot segment. Thus the actual sample division among segments varied among sectors. Only the *Energy Sector* lacked data in all segments, however. Only a 12-knot function was estimated for that sector.

¹² The coverage ratio will generally be a bit less than unity because of losses in the generation and transmission of electricity. Note that in the presence of fixed country effects, only changes of the coverage ratio over time within countries has implications for estimated income (and time) effects.

¹³ This technique is closely related to the use of the “coverage ratio” and “primary product specialization ratio” to deal with aggregation and measurement problems in the empirical study of industrial organization issues using data from the US Census of Manufactures.

General statements about 24-knot functions in what follows should be understood to mean 12-knot functions in the context of the *Energy Sector*. Simplifications to nested 12- or 6-knot spline functions were always strongly rejected on statistical grounds, though, as we shall argue, 6-knot functions generally provide good summary descriptions of the data.

3. ESTIMATION RESULTS

Table 2 shows the percentage of variance explained in the various sectors and sub-sectors using variants of equation (1). The coefficient of R_{it} , ρ , was always highly significant for the three main sectors and for *Households*; it never approached significance for *Agriculture*, *Energy Sector*, or *Non-Energy Uses*. For comparison purposes, we also include results from fitting the corresponding models (without the coverage ratio term, of course) to the 2900 available observations on *Total Apparent Consumption*.¹⁴ In general, the results from using this dependent variable are broadly consistent with those obtained earlier (with a different sample) in SSJ. Table 3 provides estimated coefficients of 6-knot spline income effect functions in specifications also including time and country fixed effects, along with related statistics.

Country Fixed Effects

Table 2 shows that country effects alone explain about 96.5 percent of the variance in total apparent consumption, consistent with our earlier aggregate analysis.

Country effects explain slightly more of the variances in sector and sub-sector

¹⁴ We did not drop any data on total apparent consumption because of suspicious year-to-year movements. Of the 3124 potential observations on that quantity in our 142-country database, 218 were dropped because of lack of GDP data, and 6 were dropped because of lack of energy consumption data.

demands. This difference is consistent with the existence of significant international differences in the structures of the broad sectors we consider.

In general, OECD nations, oil exporters, and countries with centrally planned economies tend to have large estimated country fixed effects. Countries with low estimated α_i tend to be poor countries where real GDP measurement is relatively difficult and where the traditional fuels we have excluded from our data are relatively more important. For the 109 countries for which we have estimated α_i , corresponding to all of the first four columns in Table 2, the cross-section correlations between the α_i and Y_{it} for 1985 are between 0.72 (for *Industry and Construction*) and 0.88 (for *Households and Others*).

As these results suggest, the estimated country fixed effects for the various sectors are highly but not perfectly correlated. Focusing on the first four columns in Table 2, these correlations range from 0.72 (*Transportation with Industry and Construction*) to 0.92 (*Total Apparent Consumption with both Industry and Construction and Households and Others*).

Using the 6-knot income-effect specification, the estimated α for the United States ranks fourth for *Total Apparent Consumption*, first for *Transportation*, second for *Households and Others*, but sixteenth for *Industry and Construction*. This is consistent with the argument that US energy consumption is relatively high, even controlling for per-capita income, primarily because Americans tend to drive a lot and to live in large, air conditioned homes—not because US industry is particularly energy-inefficient.

Time Fixed Effects

Despite substantial rises and falls in world oil prices since 1970, Table 2 shows that time effects alone have little ability to explain either aggregate or sectoral energy consumption. Figure 1 shows the estimated time effects corresponding to the 6-knot estimates of equation (1) for the first four columns in Table 2.¹⁵ The estimates for the three sectors are not highly correlated with each other or with real oil prices on the world market. Correlations with the estimated time effects for *Total Apparent Consumption* exceed 0.80 only for *Households and Others* and *Households*. Apart from *Households and Others* and its sub-sectors, no pairwise correlation between estimated time effects exceeds 0.80, and many are negative.

The estimated time effects clearly embody more than world oil prices and almost certainly reflect more than (unobserved) domestic energy prices. We suspect that technical change -- particularly the diffusion of energy-using technologies -- is playing an important role here. The rapid growth in the time effects for *Households and Others* and its sub-sectors is at least suggestive of a pattern in which new energy-using household appliances (including televisions, microwaves, air conditioners, and others) and agricultural technologies (including tractors and harvesting machines) are developed and become more widely available, even as energy efficiency in industry and in the energy sector increase.¹⁶

Income Effects

Table 2 shows that per-capita GDP alone explains over 80 percent of the variance of log per-capita energy consumption overall and in the three major sectors.

¹⁵ Estimates corresponding to 24-knot functions are quite similar.

The lower predictive power of domestic income in *Agriculture* and the *Energy Sector* likely reflects differences in national endowments and the growing importance of international markets in these areas. In particular, the high income coefficients estimated for the *Energy Sector* likely reflect the fact that countries in which the petroleum sector is large relative to the national economy because of favorable resource endowments also tend to have high per-capita incomes.

Table 2 also shows that time and income effects explain substantial fractions of within-country variations over time in energy demands. Finally, a comparison of the statistics for 6-knot and 24-knot specifications indicates that they have nearly identical explanatory power—though, as noted above, because sample sizes are so large, F-tests overwhelmingly reject the restrictions imposed in going from 24-knot specifications to 12-knot or 6-knot simplifications.

Figures 2-5 help one to understand the main patterns in the data.¹⁷ Figure 2a, for instance, plots residuals from the 24-knot specification for *Total Apparent Consumption* as deviations from the corresponding income/consumption relation, plotted using country and time effects corresponding to the US in 1970. This plot makes it clear that the snake-like estimated income/consumption relation tracks the data well and that the estimates of the highest and lowest segments may be affected by a few observations with extreme values of income. One of the cardinal virtues of the spline approach in this context is that influence of observations of this sort can be

¹⁶ The short-run irreversibility of investments energy-using machinery and appliances suggests the possibility of hysteresis effects, which we have begun to explore in other work.

¹⁷ It may be useful in interpreting these Figures to note that 6 = ln(403), 7 = ln(1097), 8 = ln(2981), 9 = ln(8103), and 10 = ln(22026). Plots in levels rather than logs seem to convey less information.

confined to a few estimated elasticities, a virtue not shared by polynomial specifications.

Figure 2b shows that the 6-knot income effect specification for *Total Apparent Consumption* tracks the 24-knot function well overall. Coefficients for the 6-knot specification are shown in Table 3. Both 6-knot and 24-knot specifications show relatively high elasticities at income levels corresponding to the second and third 6-knot segments, along with reductions in demand responsiveness as income grows beyond this range. This is broadly consistent with a number of previous studies that have also found that the economy-wide income elasticity of energy demand falls with income: see Bates and Moore (1992), Ebohon (1996), and the references he cites.

Note, though, that the 6-knot specification shows a negative estimated income elasticity at the highest income levels, and Table 3 indicates that this estimate differs significantly from zero, while Figure 2b shows a near-zero estimate at the highest income levels for the 24-segment specification. In fact, in the 24-segment specification, the estimated income elasticity is negative and significant only for the third segment from the top, corresponding to incomes of between \$10,435 and \$11,756 1985 US dollars. This is surely more likely to reflect some sort of isolated measurement problem than a real economic phenomenon. We are on balance fairly confident that beyond per-capita incomes of \$1,500 or so (see Table 3), there is a tendency of the economy-wide income elasticity of demand for energy to fall with per-capita income, but the evidence for a negative elasticity at high income levels is, in this sample, less than compelling.¹⁸

¹⁸ As discussed above, the sample here excludes some observations retained in the SSJ sample. SSJ found the negative relation at high incomes stronger for carbon dioxide emissions than for energy consumption. As we

Figures 3a and 3b indicate clearly that the 24-segment *Industry and Construction* estimates for the highest and, especially, the lowest income levels reflect outliers in the data. The 6-knot specification seems more reasonable at the low end. In this case, none of the three specifications yields a significant negative high-income elasticity. The income elasticity of demand in this sector appears to rise toward unity at income levels a bit above those at which this occurs for *Total Apparent Consumption*, and there is clear evidence of a drop in elasticity (though not below zero) only at the highest income levels. Note also that the estimated time effects for this sector trend downward after 1973, implying, on average, a fall in per-capita energy consumption in this sector over time.

The *Transportation* sector, illustrated in Figures 4a and 4b, again shows odd behavior at low income levels. The income elasticity at the lowest income levels is clearly dominated by a few outliers, but, as the 6-knot results indicate, the odd behavior extends fairly far up the income range. The 24-knot income effect is particularly snake-like, as it tracks odd clusters of data points, and the 6-knot function, which effectively ignores the few outliers at high income levels, seems much more reasonable. The latter function, like the functions reported in the first two columns in Table 3, shows a reduction in the income elasticity of energy demand at high income levels. For *Transportation*, however, this reduction seems to occur at a per-capita income level of around \$4,000 1985 US dollars and to be relatively small: from about 0.75 to about 0.50. There is essentially nothing here but a few outliers to suggest a negative elasticity

note there, this likely reflects a general tendency of rich nations to shift away from coal and to natural gas for environmental reasons.

in the transportation sector at high income levels. Energy use in transportation appears to rise steadily with per capita income.

Figure 5a and 5b show, once again, anomalous behavior at low income levels in the *Households and Others* sector. As Table 3 suggests, this seems mainly to reflect data on *Households*. (The corresponding graphs for *Agriculture* show that its very large estimated lowest-income elasticity is entirely the product of a few observations for Ethiopia and Malawi.) The Figures suggest a relatively constant income elasticity of around 0.5 until a per-capita income of about \$9,000, after which the elasticity appears to go negative -- particularly in the *Households* sub-sector. Note also that, as noted above, estimated average time trends are unusually large for both *Households and Others* and for *Households*. If the 6-knot specification is estimated without time effects, all estimated top segment income elasticities are positive, though none exceed 0.25.

Finally, a few words on the estimates for the *Energy Sector* and *Non-Energy Uses* are in order. The large negative lowest segment elasticity for the *Energy Sector* shown in Table 3 is driven by a few observations from Indonesia, Kenya, India, and, especially, Ethiopia. Since demand for oil, at least, is determined on world rather than domestic markets, it is not clear how seriously to take the estimated relation between *Energy Sector* energy demand and domestic GDP at higher income levels. Moreover, as noted above, large positive elasticities may simply reflect the fact that countries with large oil reserves tend to have both large energy sectors and relatively high per-capita incomes.

The lowest-segment elasticity for *Non-Energy Uses* mainly reflects odd observations for Ethiopia and Malawi. The negative elasticity reported in Table 3 corresponds to an apparent regime change, highly significant in the 24-knot specification and highly visible in graphs corresponding to Figures 2a-5a, at a per-capita income of about \$2,000 1985 US dollars. Above that level, the usual pattern of a declining income elasticity is apparent.

Sectoral Shares

Figure 6 provides a convenient visual summary of the implications of our analysis for the relation between economic development and the structure of energy consumption. For all major sectors except *Energy*, the figure shows the ratio of predicted consumption in that sector to the total predicted consumption of the included sectors. The *Energy Sector* was excluded from this figure because, for several reasons discussed above, its estimated Engel curve seems unlikely to reflect a general relation between economic development and energy demand.¹⁹ The lowest income class in Table 3 has been excluded because our general concerns about data quality for the poorest countries have been confirmed by the anomalous elasticity estimates reported in Table 3.

Predictions were calculated using the fixed effects for the U.S. in 1970. This choice only affects the levels of the sector-specific curves—via the relative sizes of sector-specific fixed effects—not their shapes. Since country fixed effects are highly

¹⁹ When the *Energy Sector* is included, its predicted share over the income interval shown in Figure 6 rises from about three percent when per-capita GDP is \$1,000 to about 10 percent at \$10,000. The predicted share of the Energy Sector then doubles as per-capita GDP rises to \$30,000. This last segment seem likely to reflect the experience of a few oil-exporting nations.

correlated across sectors, one would not expect using another nation's fixed effects to alter the levels of these curves substantially. Because time effects diverge over time (see Figure 1), use of another base year would generally have more impact. In particular, use of a later year would tend to raise the *Households & Other* curve relative to the others because its time effects are estimated to grow most rapidly.

Figure 6 shows a steady decline in the predicted share of energy consumption accounted for by this sector as income rises. A graph of predicted shares within the *Households & Other* sector shows relatively stable shares for *Agriculture* (between 2.5 and 4.5 percent) and for *Households* (between 55 and 65 percent). The rise in *Transportation's* predicted share, particularly in the highest income range, is consistent with both conventional wisdom and other studies; see, for instance, Dunkerley and Hoch (1987) on developing countries. Finally, we are unaware of any study that finds the inverse-U pattern that Figure 6 reveals for the share of *Industry & Construction*, though Howarth and Schipper (1991) and others have pointed to declines in industrial energy-intensity in OECD countries.

4. CONCLUDING REMARKS

We believe this study demonstrates both that the U.N. sectoral consumption data contains interesting information and that extracting that information requires considerable care. Our estimated reduced-form Engel relations for major economic sectors are not simply replicates of the aggregate relation, and for the most part the differences among our estimated sectoral relations are plausible. Only our income

elasticity estimates for very low income levels seem primarily driven by measurement error.

We have found substantial differences among major sectors in patterns of country, time, and income effects. As at the aggregate level, both country and income effects explain a large fraction of variation in sectoral energy demands, while time effects have considerably less explanatory power. Estimated country fixed effects are highly correlated across sectors, while time effects are not. At the very least, this latter finding makes it clear that time effects here reflect more than changes in energy prices on world markets. It seems likely that differences in estimated time effects reflect differences in sector-specific directions of technological change rather than fuel-specific or sector-specific differences in domestic price trends.

Differences in patterns of income elasticities are of particular interest in connection with the process of economic development and with forecasting future energy demands and carbon dioxide emissions. There is general evidence that income elasticities decline with income, particularly at the highest income levels. The negative top-segment elasticity that we found in SSJ appears to be driven entirely by the *Households & Other* sector. As per-capita income rises, our estimates imply that this sector's share of aggregate energy consumption tends to fall, while the share of *Transportation* tends to rise, and the share of *Industry & Construction* follows an inverse-U pattern. Absent a dramatic change in motor vehicle technology, this implies a continuing rise in the relative importance of petroleum as an energy source.

References

- Alcamo, J. et al (1995). "An Evaluation of the IPCC IS92 Emission Scenarios." Ch. 6 in J.T. Houghton et al, eds., *Climate Change 1994*. Cambridge: Cambridge University Press.
- Bates, Robin W. and Edwin A. Moore (1992). "Commercial Energy Efficiency and the Environment," Policy Research Working Paper WPS 972. Washington: The World Bank, Environmental Policy and Research Division, September.
- Bhatia, Ramesch (1987). "Energy Demand Analysis in Developing Countries: A Review." *Energy Journal*, 8: 1-33.
- Dahl, Carol (1992). "A Survey of Energy Demand Elasticities for the Developing World." *Journal of Energy and Development*, 18 (Autumn): 1-47.
- Dunkerley, Joy and Irving Hoch (1987). "Energy for Transport in Developing Countries." *Energy Journal*, 8: 57-72.
- Ebohon, Obas John (1996). "Energy, Economic Growth and Causality in Developing Countries." *Energy Policy*, 24 (5): 447-53.
- Howarth, Richard B. and Lee Schipper (1991). "Manufacturing Energy Use in Eight OECD Countries: Trends through 1998." *Energy Journal*, 12: 15-40.
- IPCC (1996). *Climate Change 1995: The Science of Climate Change*. Contribution of Working Group I to the Second Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge: Cambridge University Press.
- Marland, Gregg et al (1989). *Estimates of CO₂ Emissions from Fossil Fuel Burning and Cement Manufacturing, Based on the United Nations Energy Statistics and the U.S. Bureau of Mines Cement Manufacturing Data*, ORNL/CDIAC-25, NDP-030. Oak Ridge, TN: U.S. Department of Energy, Oak Ridge National Laboratory, May.
- Pindyck, Robert S. (1979). *The Structure of World Energy Demand*. Cambridge, MA: MIT Press.
- Schmalensee, Richard, Thomas M. Stoker, and Ruth A. Judson (1998). "World Carbon Dioxide Emissions: 1950 -- 2050," *Review of Economics and Statistics*, forthcoming.
- Summers, Robert and Alan Heston (1991). "The Penn Mark IV Data Table," *Quarterly Journal of Economics*, 106 (May), 327-368.

United Nations (1994a). *1992 Energy Balances and Electricity Profiles*. New York: United Nations.

United Nations (1994a). "The 1992 United Nations Energy Statistics Database." Mimeo, United Nations, May.

United Nations (1995). *1993 Energy Statistics Yearbook*. New York: United Nations.

Figure 2a: Total Apparent Consumption

Figure 2b: Total Apparent Consumption

Figure 3a: Industry and Construction

Figure 3b: Industry and Construction

Figure 4a: Transportation

Figure 4b: Transportation

Figure 5a: Households and Other

Figure 5b: Households and Other

Table 1: Summary of Sectoral Energy Consumption Data*

Sector	Total Sample	OECD Sample	Countries	Population Coverage	Per Capita Mean	Per Capita Coeff. of Var.	Per Dollar Mean	Per Dollar Coeff. of Var.
Industry & Construction	1761	413	119	69.97	17.10	1.96	2.63	1.33
Transportation	1832	434	119	68.82	7.69	1.31	1.26	0.70
Households & Other	1837	420	123	69.85	12.10	1.42	1.77	1.03
Households	1921	446	121	65.57	7.41	1.50	1.08	1.15
Agriculture	1096	335	79	59.11	1.06	1.21	0.19	1.06
Energy Sector	1044	407	69	59.00	5.08	1.72	0.73	1.58
Non-Energy Uses	1286	383	92	56.76	2.77	1.22	0.40	1.14

* “Total Sample,” “OECD Sample,” and “Countries” give the total number of observations, the number of observations on OECD nations, and the total number of countries, respectively, in the sample for the sector indicated. “Population Coverage” gives the percentage of the world population in 1980 (from the World Bank’s *World Development Report*) accounted by our samples for that year. “Per Capita Mean” gives the mean values of consumption measured in 10^{12} Btus per capita, and the column that follows gives the corresponding coefficients of variation. “Per GDP Mean” gives the mean values of consumption measured in 10^{12} Btus per thousand 1985 dollars of GDP, and the column that follows gives the corresponding coefficients of variation.

Table 2: Percentages of Variance Explained*

Effects	Total Apparent Consumption	Industry & Construction	Transportation	Households & Other	Households	Agriculture	Energy Sector	Non-Energy Uses
<i>Total Variance:</i>								
Country Only	96.47	98.26	97.75	98.46	98.26	96.53	96.78	97.07
Time Only	0.01	2.04	0.06	1.10	1.57	1.76	2.15	0.74
Income Only (24)	82.43	80.93	88.17	84.64	81.18	53.08	58.58	77.78
Income Only (6)	82.18	79.93	87.56	83.35	80.17	49.30	54.96	76.41
All Effects (24)	97.73	98.81	98.46	99.05	98.81	97.60	97.43	97.96
All Effects (6)	97.66	98.75	98.31	99.01	98.74	97.48	97.31	97.75
<i>Within-Country Variance:</i>								
Income Only (24)	33.64	30.52	31.06	27.85	27.02	23.36	14.98	28.47
Income Only (6)	31.63	27.06	23.88	25.14	23.06	18.77	12.31	21.27
All Effects (24)	35.63	31.64	31.67	38.27	31.71	30.99	20.25	30.29
All Effects (6)	33.70	28.35	24.66	35.84	27.89	27.38	16.69	23.13

* Figures in parentheses are numbers of knots/segments in the income effect spline function -- except that for the Energy Sector, "(24)" denotes results from 12-knot spline functions, as discussed in the text. Except for Total Apparent Consumption, the income effects include R, the log of the coverage ratio.

Table 3: Estimated Coefficients of 6-Knot Spline Income Effect Functions*

Income Range: 1985\$/capita	Total Apparent Consumption	Industry & Construction	Transportation	Households & Other	Households	Agriculture	Energy Sector	Non-Energy Uses
≤ 823	0.219 (0.096)	0.287 (0.157)	-0.187 (0.117)	-0.464 (0.116)	-0.699 (0.122)	1.749 (0.640)	-5.094 (1.608)	-0.107 (0.640)
823 - 1,430	6.40 1.098 (0.081)	1.53 0.618 (0.117)	5.28 0.704 (0.102)	6.66 0.613 (0.092)	8.80 0.817 (0.100)	-2.24 0.148 (0.212)	3.76 1.599 (0.285)	2.36 1.522 (0.182)
1,430 - 2,545	2.43 1.400 (0.082)	2.39 1.046 (0.117)	1.22 0.892 (0.103)	-2.45 0.229 (0.112)	-5.68 -0.104 (0.113)	1.33 0.516 (0.175)	-0.43 1.410 (0.270)	-6.44 -0.201 (0.180)
2,545 - 4,249	-4.83 0.784 (0.085)	-0.80 0.899 (0.121)	-0.74 0.774 (0.105)	1.96 0.538 (0.096)	5.64 0.790 (0.104)	2.51 1.189 (0.170)	-1.03 1.072 (0.176)	8.34 1.899 (0.160)
4,249 - 8,759	-3.16 0.394 (0.075)	-0.10 0.882 (0.088)	-2.16 0.466 (0.079)	-0.30 0.500 (0.074)	-1.46 0.582 (0.082)	-4.02 0.203 (0.148)	-1.73 0.684 (0.123)	-5.37 0.744 (0.119)
≥ 8,759	-5.98 -0.312 (0.085)	-8.08 -0.186 (0.097)	0.47 0.521 (0.083)	-7.46 -0.349 (0.086)	-7.83 -0.396 (0.093)	-4.34 -0.767 (0.198)	0.83 0.837 (0.157)	-2.58 0.293 (0.134)
ρ	--	0.435 (0.036)	0.177 (0.030)	0.174 (0.027)	0.191 (0.027)	0.016 (0.060)	0.023 (0.047)	-0.020 (0.026)
Trend	0.69	-0.36	0.17	1.83	1.28	2.23	-1.03	--

* Figures in parentheses are standard errors. Income coefficients are elasticities of per-capita consumption with respect to per-capita income. Figures between elasticity estimates are the t-statistics for the null hypotheses that the adjacent elasticities are equal. “ρ” is the coefficient of R, the log of the sectoral coverage ratio, and “Trend” is the estimated slope, if significant at 5%, in a least-squares regression of estimated time effects on a trend variable, expressed as an annual percentage increase.

Figure 1: Estimated Time Effects, 6-Knot Spline Specification

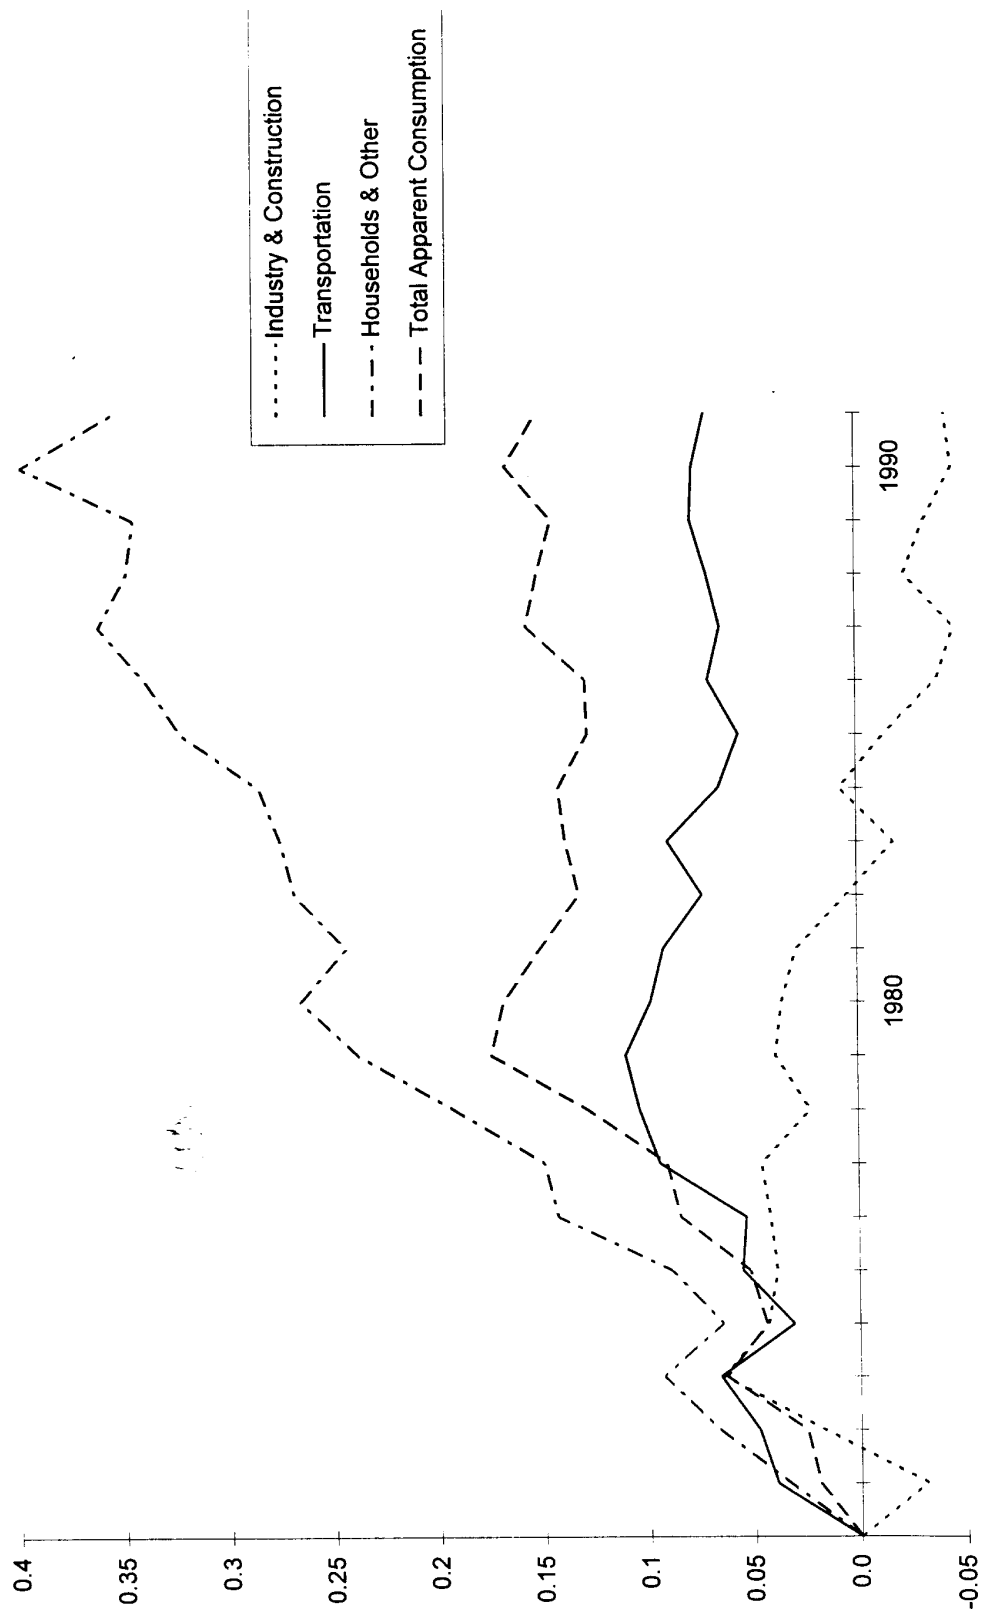


Figure 2a: Total Apparent Consumption

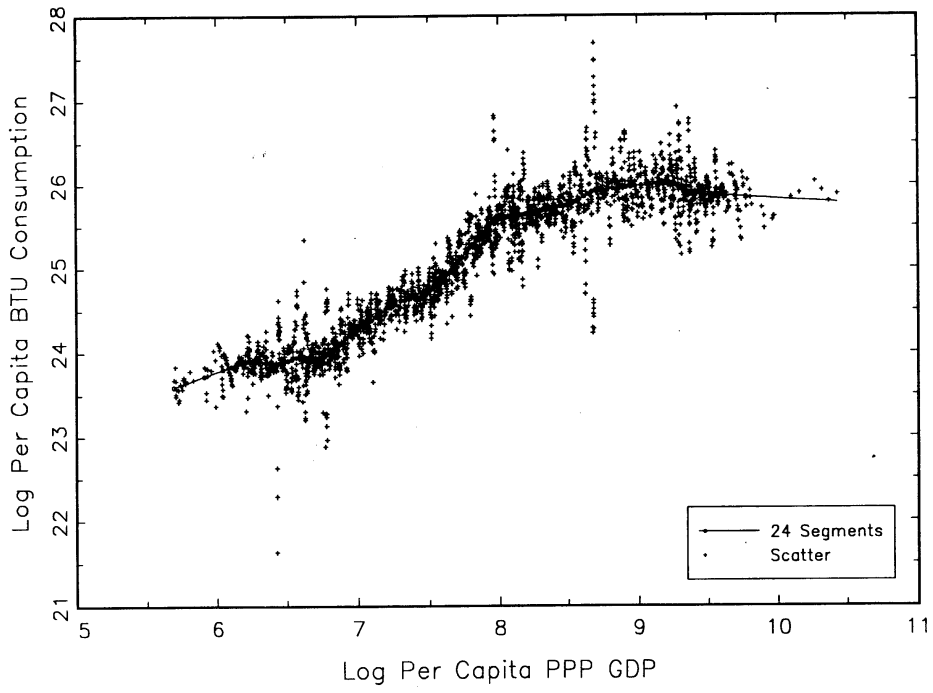


Figure 2b: Total Apparent Consumption

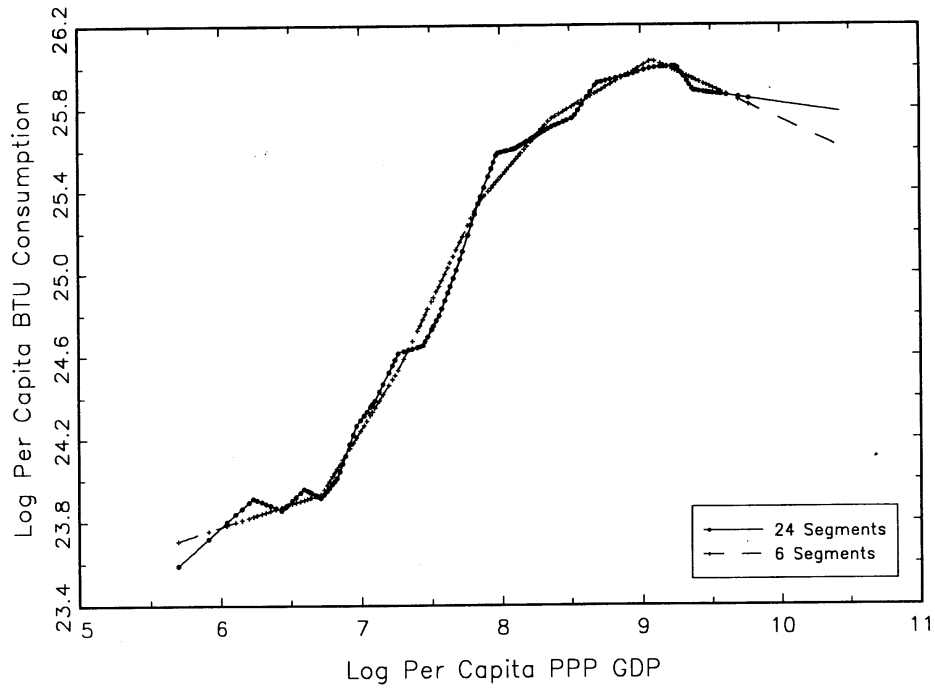


Figure 3a: Industry and Construction

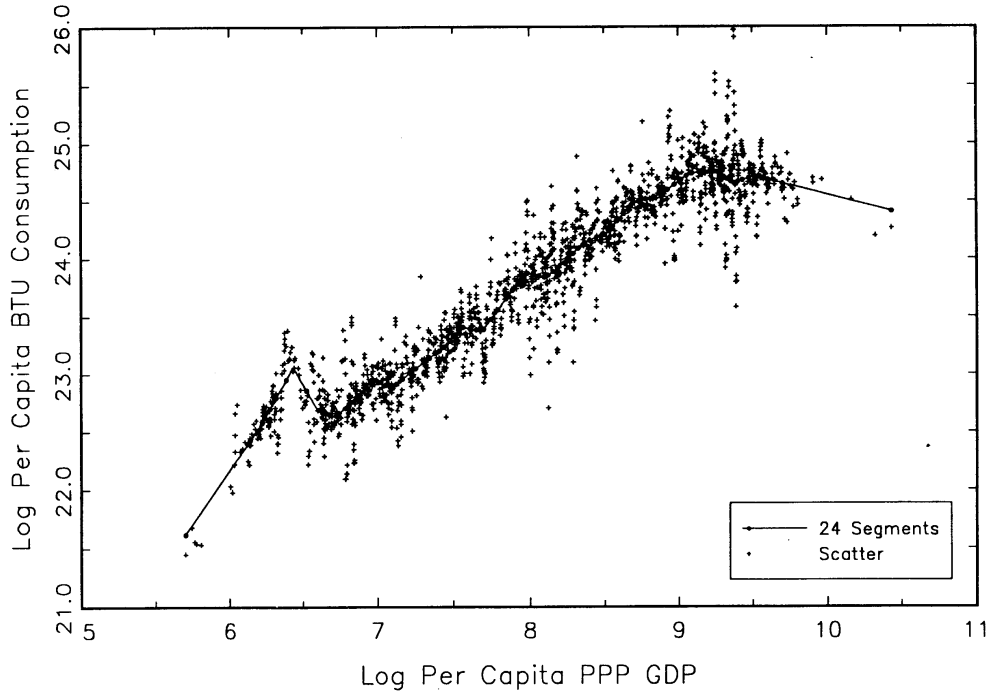


Figure 3b: Industry and Construction

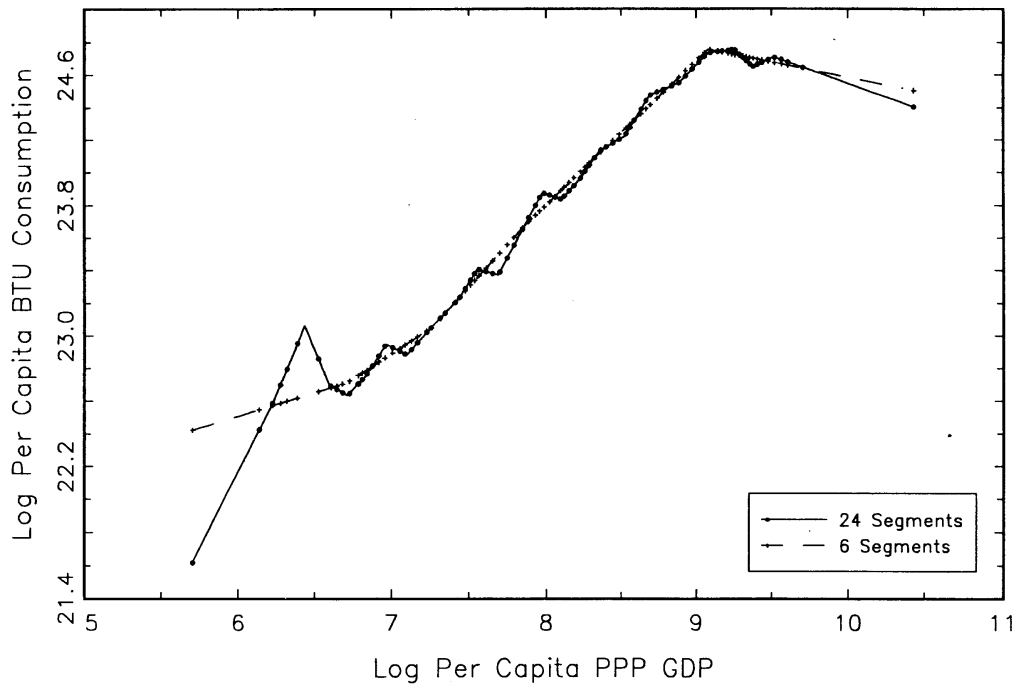


Figure 4a: Transportation

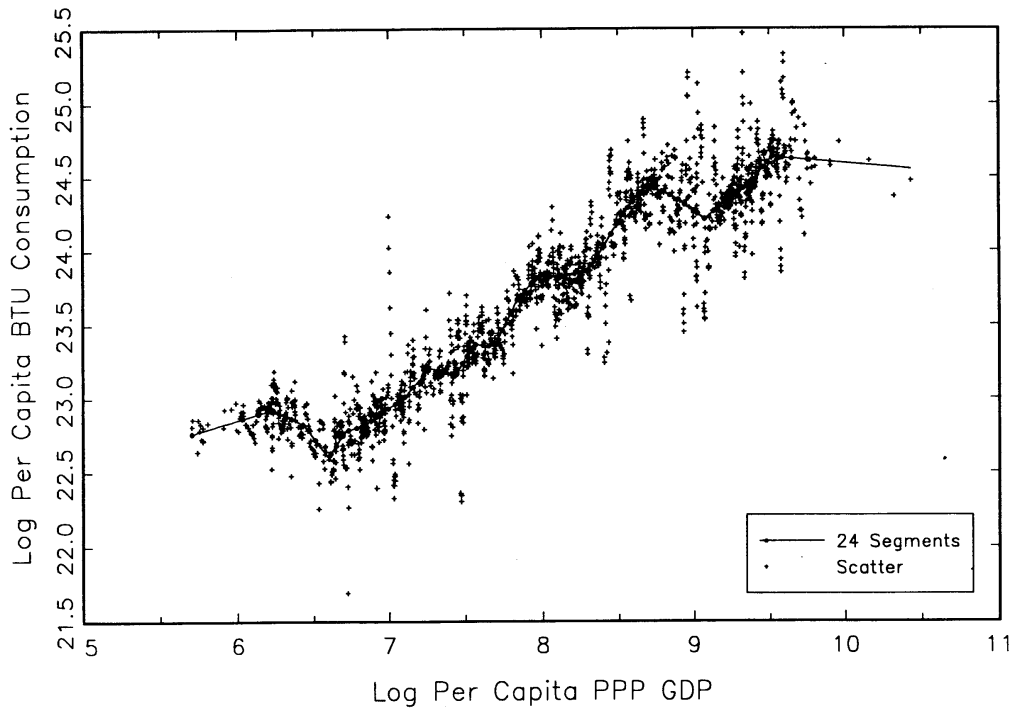


Figure 4b: Transportation

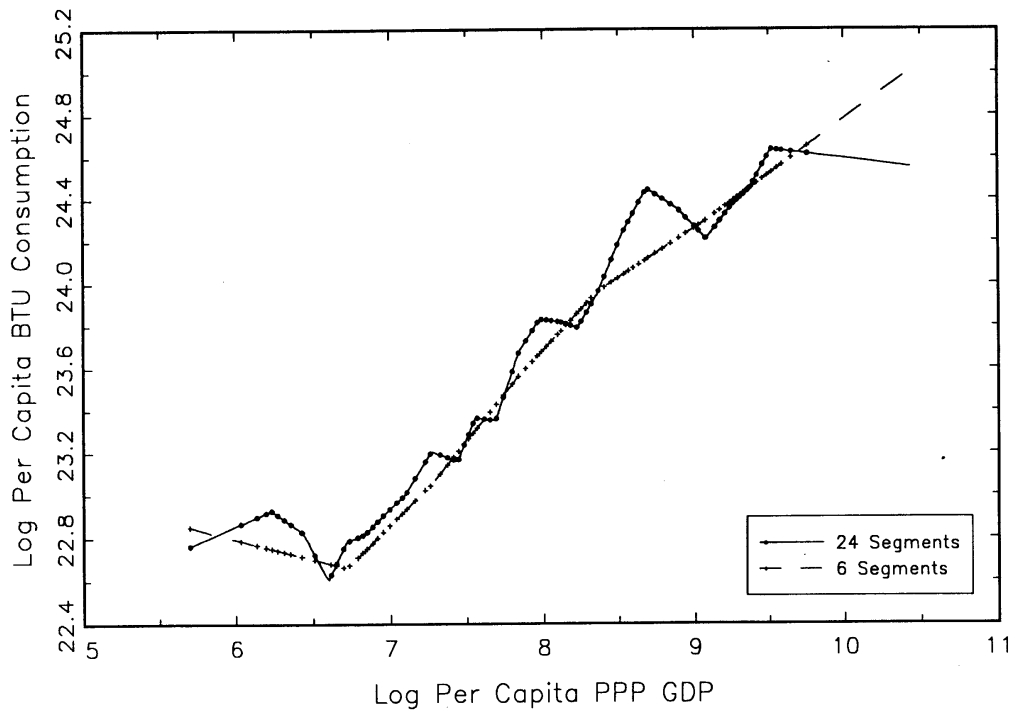


Figure 5a: Households and Other

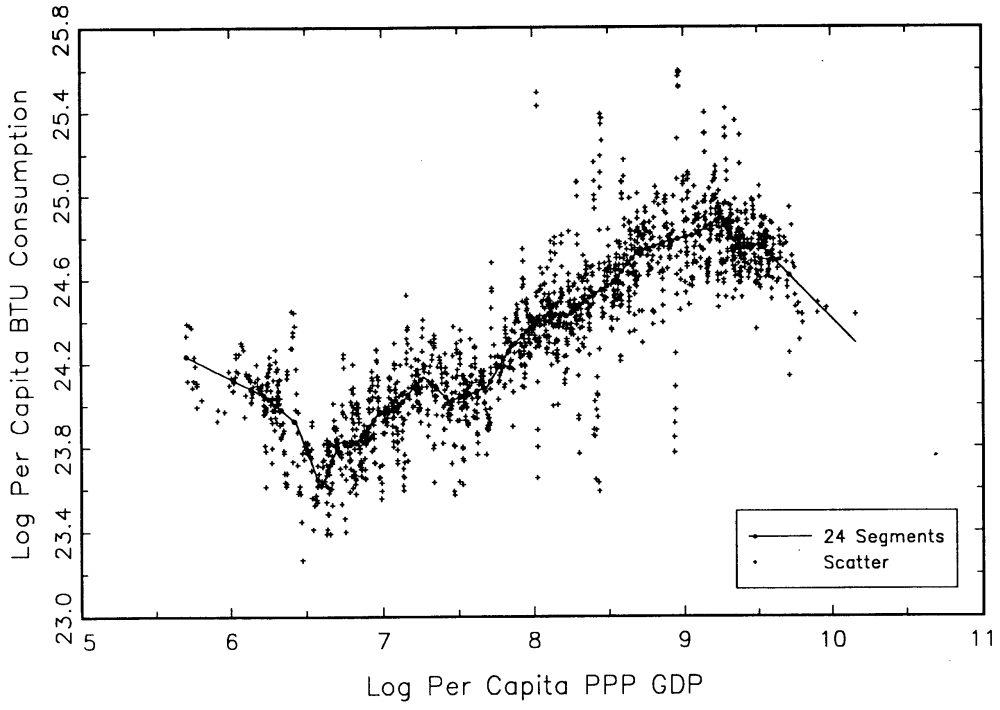
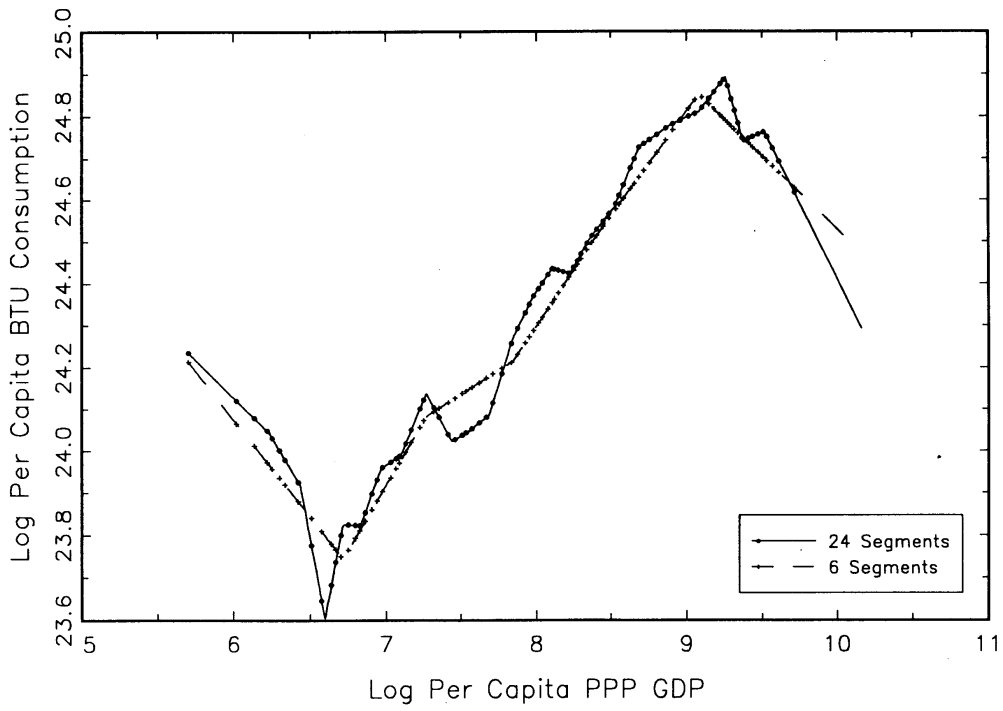


Figure 5b: Households and Other



**Figure 6: Sector Shares and Economic Development,
Energy Sector Excluded**

