

Forecasting Regional Long-Run Energy Demand: A Functional Coefficient Panel Approach*

Yoosoon Chang[†], Yongok Choi[‡], Chang Sik Kim[§], J. Isaac Miller[¶], Joon Y. Park^{||}

Abstract

Previous authors have pointed out that energy consumption changes both over time and nonlinearly with income level. Recent methodological advances using functional coefficients allow panel models to capture these features succinctly. In order to forecast a functional coefficient out-of-sample, we use functional principal components analysis (FPCA), reducing the problem of forecasting a surface to a much easier problem of forecasting a small number of smoothly varying time series. Using a panel of 180 countries with data since 1971, we forecast energy consumption to 2035 for Germany, Italy, the US, Brazil, China, and India.

This Version: November 25, 2019

JEL Classification: C14, C23, C51, Q43

Key words and phrases: functional coefficient panel model, functional principal component analysis, energy consumption

*The authors are grateful to participants of the 2018 International Conference on the Economics of Oil, Fundação Getulio Vargas, and the 2019 Workshop on Energy Economics: Econometric Analysis of Energy Demand and Climate Change, Sungkyunkwan University and Korea Power Exchange, for useful feedback.

[†]Department of Economics, Indiana University

[‡]School of Economics, Chung-Ang University

[§]Department of Economics, Sungkyunkwan University

[¶]Department of Economics, University of Missouri

^{||}Department of Economics, Indiana University and Sungkyunkwan University

1 Introduction

Long-term forecasts of local/national/international energy consumption are important components not only in the domain of energy policy, but also those of policies aimed at economic growth, economic diversification, and economic inequality, climate change, agriculture, and other areas. After all, energy is one of the main factors of production of economic output, and its consumption is an inescapable reality in modern economies.

As does all consumption, energy consumption changes with wealth, and this relationship is generally considered to be nonlinear. Such nonlinearity may result from the so-called environmental Kuznets curve, or its origin may be a more subtle shift in the sectoral shares of an economy, moving from low-intensity agriculture to high-intensity industry to low-intensity services. Some authors have modeled the implied inverted U-shape parametrically, using a quadratic function (Galli, 1998; Medlock and Soligo, 2001; and Richmond and Kaufmann; 2006a,b), while others have used more flexible nonparametric or semiparametric methods (Judson *et al.*, 1999; Luzzati and Orsini, 2009; and Nguyen-Van, 2010).

At the same time, the relationship between income and energy consumption changes over time even for a fixed income level. For example, the oil price shocks of the 1970's led a drive to improve technical efficiency in multiple economic sectors in oil-importing countries. Such efficiency naturally decreased the impact of income changes on energy consumption. As another example, consider the proliferation of electricity-hungry electronic devices over the past few decades. The effect of this proliferation has likely been the reverse: a modern consumer likely uses more electricity to drive personal computing and telecommunications than his/her precursor, say, three decades ago.

The present work continues a series of analyses by the authors (Chang *et al.*, 2016a, 2019b) that model energy consumption as a nonparametrically specified and time-varying

function of income using a panel of countries, effectively addressing both the nonlinearity and instability over time. The fact that we use data by country results from data availability, but the methodology is appropriate for any cross-sectional unit: country, state or province, regional transmission organization or independent system operator for electricity consumption, etc.

To the best of the authors’ knowledge, ours is the first analysis to address out-of-sample forecasting of energy consumption using a nonparametrically specified, time-varying panel model. Forecasting the surface estimated by nonparametric methods is accomplished by reducing the problem of forecasting along a continuum to the simpler problem of forecasting the leading principal components identified using functional principal components analysis (FPCA, see Hörmann and Kokoszka, 2012).

We are certainly not the first authors to use FPCA in forecasting. Pioneering work by Bosq (2000) and Besse *et al.* (2000) use functional autoregressive models for out-of-sample predictions, and others have employed similar methods (Hyndman and Ullah, 2007; Aue *et al.*, 2014, *inter alia*). As shown by Chang *et al.* (2019b), using panel data allows us to decompose the relationship between income and consumption into heterogeneous and common components, and it is the latter to which we apply FPCA.

The common components of the wealthier countries show increasing energy efficiency (declining energy intensity)¹ similar to that documented for the US and other developed countries (Kaufmann, 2004; Webster *et al.*, 2008; US Energy Information Administration, 2013; Csereklyei *et al.*, 2016). Developing countries generally show declining energy efficiency

¹We use the term “energy efficiency” in the sense that an increase in energy efficiency (an “autonomous energy efficiency increase” in the language of Kaufmann, 2004) correlates with a decrease in energy intensity. Importantly, energy efficiency may have little to do with technical efficiency and more to do with the sectoral composition of a region. As Bradford (2018) points out, Iceland has among the highest energy intensities in the world due to its abundant hydroelectric and geothermal resources, but not because its heavy metals industry is technically inefficient.

(increasing energy intensity), but a developing country may “graduate” from that group if sectoral shifts or economic diversification increases energy efficiency.

This research relates to three presentations given at the conferences in Brazil and South Korea around which the special issue is organized. One of the present authors, Choi, discussed the grouping and estimation of Chang *et al.* (2019b) at the venue in Brazil in 2018. Another, Miller, discussed the forecasting methodology and results at the same conference. A third, Park, discussed application of the methodology to the electricity demand at the venue in South Korea in 2019. Indeed, it was the problem of forecasting South Korean electricity demand that spurred the development of the concepts and methodologies in these two papers, and for which these models are in fact employed.

The remainder of the paper is structured as follows. In Section 2, we briefly introduce the forecasting model and methodology, although we defer to previous work to explain details of estimation. Section 3 contains the empirical results and we conclude with Section 4. An appendix contains details of the bootstrap methodology we utilize to obtain forecast intervals.

2 Forecasting Model and Methodology

In forecasting energy demand, we employ a panel model given by

$$y_{it} = \alpha_i + (\gamma_i + \beta(t, x_{it}))x_{it} + \varepsilon_{it} \quad (1)$$

for region $i = 1, \dots, N$ and year $t = 1, \dots, T$. The data are y_{it} , log real final energy consumption per capita (log tonne of oil equivalent per capita), and x_{it} , log real GDP per capita (log 2011 dollars per capita in output-side chained PPPs from Penn World Tables 9.0), for 180 countries

over 1971-2015. For expositional simplicity, we will refer to these variables simply as demand or consumption and income, respectively. Intercepts α_i are heterogeneous, and slopes have both a heterogeneous component γ_i and a common component $\beta(r, x)$, which may change over time and over income, but is not country-specific.

In a related paper (Chang *et al.*, 2019b), we motivate the model in equation (1) from a more general panel framework and argue for a structural interpretation of the parameters in the context of world energy consumption. In the present context of forecasting, the structural interpretation is not necessary.

Semiparametric estimation of the panel model proceeds as described by Chang *et al.* (2019b), which closely relates to that of Chang *et al.* (2016a) adapted from Fan and Huang (2005). Readers interested in implementing the estimation procedure are referred to these papers for details. We denote estimates of these parameters by $\hat{\alpha}_i$, $\hat{\gamma}_i$, and $\hat{\beta}(r, x)$, and note that the latter estimates a surface.

Constructing a conditional forecast using a linear fixed-coefficient forecasting model is straightforward once the parameters are estimated. However, in the present context, even if we have $x_{i,T+h}$ for $h > 0$, it is not obvious what would be $\beta(T + h, x_{i,T+h})$, because the surface $\beta(r, x)$ is not estimated precisely outside of the support of the historical data.

In order to overcome this obstacle, we view (β_t) , $\beta_t \equiv \beta(t, \cdot)$, as a functional time series – i.e., a time series of functional observations – and forecast β_{T+h} for any h in the forecasting horizon. To do this, we use the functional factor model developed in Chang *et al.* (2019a). Following them, we consider a functional factor model given by

$$\delta_t = \beta_t - \bar{\beta} = \sum_{k=1}^K c_{kt} f_k + \eta_t,$$

where $\bar{\beta}$ is the temporal average of (β_t) , $f_1, \dots, f_K$ are the functional common factors with

scalar loadings $c_{1t}, \dots, c_{Kt}$, and (η_t) is the idiosyncratic functional component. The loadings on the common factors are assumed to be regular time series, whereas the idiosyncratic component is regarded as being irregular and representing a sporadic and transient time effect.

Under this setup, Chang *et al.* (2019a) obtain the least squares estimates of the functional factors and their loadings. The estimates of the functional factors $f_1, \dots, f_K$ are given by the orthonormal eigenfunctions $v_1, \dots, v_K$ of $\sum_{t=1}^T (\delta_t \otimes \delta_t)$ associated with its K -largest eigenvalues, and the estimates of their loadings $c_{1t}, \dots, c_{Kt}$ are given by $\langle v_1, \delta_t \rangle, \dots, \langle v_K, \delta_t \rangle$, where we use the notations “ $\otimes$ ” and “ $\langle \cdot, \cdot \rangle$ ” to denote the tensor and inner products in the Hilbert space of square integrable functions. Therefore, we may readily obtain the estimated functional factors $v_1, \dots, v_K$ and their loadings $\langle v_1, \delta_t \rangle, \dots, \langle v_K, \delta_t \rangle$ from functional principal component analysis (FPCA). The interested reader is referred to Hörmann and Kokoszka (2012) for a technical introduction to and survey of FPCA, Hörmann *et al.* (2015) for some extant uses of FPCA, and Bosq (2000) for fundamental theory of functional time series on which FPCA is based.

We assume that the idiosyncratic component (η_t) is dominated by the common factor component, so that we have

$$\delta_t = \beta_t - \bar{\beta} \simeq \sum_{k=1}^K \langle v_k, \delta_t \rangle v_k,$$

which may simply be regarded as the K -th order Karhunen-Loève expansion of (δ_t) using an orthonormal basis (v_k) . We show later that our assumption here is empirically justifiable.

The number K of factors can be consistently estimated by the eigenvalue ratio test of Ahn and Horenstein (2013). Note that β_t is not directly observable, so we need to use the estimate $\hat{\beta}_t = \hat{\beta}(t, \cdot)$. Using $\hat{\beta}_t$ in place of β_t , however, does not affect consistency of the test.

See Chang *et al.* (2019a) for more details.

Now we write

$$\delta_t = \beta_t - \bar{\beta} \simeq \sum_{k=1}^K \hat{c}_{kt} \hat{f}_k,$$

and forecast $\hat{c}_{1t}, \dots, \hat{c}_{Kt}$ to obtain their future values. Then we may define

$$\hat{\beta}_t = \bar{\beta} + \sum_{k=1}^K \hat{c}_{kt} \hat{f}_k,$$

from which we have our forecast

$$\hat{\beta}(t, x) = \hat{\beta}_t(x)$$

of $\beta(t, x)$ up to the forecasting horizon.

By approximating the surface by its K leading factors following Hyndman and Ullah (2007), we vastly simplify the task of forecasting a surface to that of forecasting a small number of smoothly varying time series – i.e., the factor loadings, $\hat{c}_{1t}, \dots, \hat{c}_{Kt}$. The smoothness of the estimated surface naturally depends on two bandwidth parameters, details of which are explained by Chang *et al.* (2019b). We use exactly the same bandwidth to estimate the surface for the developed group and a similar bandwidth to estimate the surface for all countries.²

Our procedure for implementing FPCA draws on that of Chang *et al.* (2016b, 2019c). We first temporally demean $\hat{\beta}_t(x)$ at each ordinate x for which we estimate the surface, labeling the result $\hat{\delta}_t(x)$ as above. Next, we conduct a wavelet decomposition of $\hat{\delta}_t(x)$. Specifically, we use $M = 1,037$ Daubechies wavelet basis functions and label the decomposition $[\hat{\delta}_t]$, which is an M -dimensional vector for fixed t .

²Specifically, $h_r = cn^{-1/6}$ and $h_x = c\tau n^{-1/6}$ with $c = 0.4$ chosen for both surfaces and $\tau = 0.261, 0.411$, the ratios of the standard deviations of t/T to those of $x_{it}/\max(x_{it})$ for the surfaces for the developed group and for all countries respectively.

Using $[\hat{\delta}_t]$, we calculate the orthonormal eigenvectors $([v_k])_{k=1}^M$ of the $M \times M$ matrix $[Q_T] = \sum_{t=1}^T [\hat{\delta}_t][\hat{\delta}_t]'$. We then approximate $[\hat{\delta}_t]$ using

$$[\hat{\delta}_t] \simeq \sum_{k=1}^K [v_k][v_k]'[\hat{\delta}_t],$$

where the eigenfunctions $([v_k])_{k=1}^K$ are the K leading eigenfunctions (factors) with the associated loadings (factor loadings) $([v_k]'[\hat{\delta}_t])$. Finally, we add back in the temporal mean $\bar{\beta}$ to approximate the estimated surface $\hat{\beta}_t$.

3 Empirical Results

We discuss our empirical results in three steps, the methodologies underlying the first two of which are explained above. In the first step, we estimate the surface $\beta(r, x)$ for points (r, x) within the span of historical time t and incomes (x_{it}) . In the second step, we decompose that surface using FPCA and extend the leading components into the future. In the third step, we forecast $y_{i,T+h}$ conditionally on $x_{i,T+h}$.

3.1 Estimating the Surfaces

Chang *et al.* (2019b) identify two sets of countries sharing a common slope $\beta(t, x_{it})$. These groups consist of countries considered to be developed and developing, roughly speaking. More specifically, they are countries whose energy efficiency patterns resemble each other. The first group, which we will call the developed group and includes Germany, Italy, and the US, shows a recent increase in energy efficiency, while the second, which we will call the developing group and includes Brazil, China, and India, does not.

For the purpose of forecasting, we estimate the surface using only the developed group




Figure 1: **Estimates of the Surface $\beta(x, r)$.** Left panel: estimates using the developed group only (46 countries, 1867 observations). Right panel: estimates using the full sample (180 countries, 7482 observations). The estimation and bias-correction procedures are described by Chang *et al.* (2019).

and that using all of the countries in our sample. Our reason for this approach is that the energy efficiency of a country already in the developed group is expected to continue to resemble that of the developed group into the future. Including the developing group in estimating this surface adds unnecessary noise to the estimation procedure. On the other hand, a country in the developing group may “graduate” from that group if sectoral shifts or economic diversification increases energy efficiency to more closely resemble that of countries in the developed group. It is therefore useful to include all countries in forecasting long-run demand for a country in the developing group.

Figure 1 shows the two surfaces estimated for the developed group (left panel) and for all countries in our sample (right panel) with incomes converted out of log scale and back into 2011 US dollars. A caveat for both surfaces is that countries’ incomes tend to increase over time. As a result, there are relatively more observations for lower incomes near the beginning of the sample and for higher incomes near the end than for higher incomes near the beginning

Country	$\hat{\alpha}_i^u$	$\hat{\alpha}_i^c$	s.e.	$\hat{\gamma}_i^u$	$\hat{\gamma}_i^c$	s.e.
Group A:						
Germany	3.149	3.005	0.519	-0.417	-0.363	0.048
Italy	-1.519	-1.221	0.333	0.006	0.019	0.020
United States	5.413	5.276	0.886	-0.570	-0.516	0.070
Group B:						
Brazil	0.445	0.375	0.544	-0.139	-0.144	0.054
China	-4.587	-4.225	0.239	0.435	0.383	0.022
India	-2.723	-2.602	0.135	0.147	0.118	0.031

Table 1: **Coefficient Estimates for Selected Countries.** Estimates for countries in the developed group employ the sample with the developed group only (46 countries, 1867 observations). Estimates for countries in the developing group employ the full sample (180 countries, 7482 observations). Superscripts u and c denote uncorrected and bias-corrected estimates, as explained in detail by Chang *et al.* (2019). The bootstrap procedure used to construct standard errors is described in the appendix.

and lower incomes near the end. We therefore expect the left-most and right-most corners of each figure to be less precisely estimated.

The surface for the developed group shows a non-monotonicity over time at all income levels, increasing for a few decades and then leveling out or decreasing. Similarly, fixing time seems to show a similar inverted U-shape over income. Roughly speaking, the least energy efficiency – or the peak of inefficiency – seems to have been achieved in middle income countries during the 1990s, when indeed energy prices were quite low providing relatively little incentive for technically efficient energy use.

In contrast, the surface estimated using all countries in the sample peaks at the highest income plotted (\$50,000) and at the most recent observation is the sample. We observe a leveling out of the surface as income increases, and it may well decline beyond \$50,000.

More telling is the contrast between the surfaces for the developed group and for all countries over time. Aside from the lowest income levels, the surface for all countries at each income level increases nearly monotonically over time, while it is clearly not monotonic for

the developed group, which seems to level out and even decline starting in the 1990s.

Tables 1 shows estimates of α_i and γ_i for selected countries, using the subsample for countries in the developed group and the full sample for countries in the developing group. The negative slope coefficients for Germany and the US suggest efficiency *gains* (being in the developed group) that increase faster with higher levels of income than those of the average of the developed group at the same level of income. The positive slope coefficient for Italy, on the other hand, suggests slower efficiency gains at high income levels. Brazil shows slower efficiency *losses* (not being in the developed group) at high income levels relative to the average of all countries, while booming China and India show faster efficiency *losses* with higher income levels relative to the average.

3.2 Extending the Surfaces

Developed Group. Extending the surfaces in Figure 1 proceeds by functional approximation. Figure 2 plots cross-sections of the surface estimated for the developed group in Figure 1 (left panel), which provides a heuristic for our approach using functional principal components. To an approximation, all of the cross-sections show a rise as income increases up to a certain point, after which the effect of income flattens dramatically. The pattern is more pronounced later in the sample.

The cumulative contributions of the three leading factors from decomposing the surface for the developed group using FPCA are 87.0%, 98.4%, and 99.4%, and the eigenvalue ratio test of Ahn and Horenstein (2013) selects two factors.

Figure 3 shows the first factor and its loading. Notice the shape of the factor: it rises as income increases up to a certain point, after which it is nearly flat. The factor loading puts more positive weight on this pattern over time. The first factor and its loading thus show patterns similar to those of the surface that they estimate.




Figure 2: **Cross-Sections for the Developed Group.** Cross-sections of the surface $\beta(r, x)$ from fixing $r = \{1971, 1980, 1990, 2000, 2015\}$.




Figure 3: **First Factor and Its Loading for the Developed Group.** Estimated using FPCA.




Figure 4: **Second Factor and Its Loading for the Developed Group.** Estimated using FPCA.




Figure 5: **Extending the Factor Loadings for the Developed Group.** The first factor loading (left) is extended along the time dimension as a constant function of time. The second factor loading (right) is extended using a logistic function of time.

The second factor and its loading show a suppressing effect (Figure 4). The factor is negative for most income levels and its loading increases strongly since the late 1990s. Their product is therefore mostly negative, suppressing the function – i.e., amplifying energy efficiency – over the last decade. Chang *et al.* (2016a) document possible explanations for a change in energy efficiency among richer countries at about this time, including a shift away from energy-intensive industries and an increasing awareness of environmental externalities of energy use following the Kyoto Protocol – i.e., from the changing slope of an environmental Kuznets curve.

The smoothness of the factor loadings allows more straightforward out-of-sample forecasting than using functional autoregression. From 2006-2015, the first factor loading is nearly flat, so we extend this loading in the time dimension by a constant function. Figure 5 (left panel) shows the result. The right panel of the figure shows an extension of the second factor loading, which is a bit more complicated. Between 1998 and 2015, the function appears to be approximately sigmoidal, so we extend it using a logistic function. Specifically, we fit $(\theta - \delta) / (1 + \exp(-\alpha(\tau - \beta))) + \delta$ over 1998-2015, where $\tau = (year - 1998)/17$, and then extend the function up to $year = 2035$. The parameter estimates (standard errors) show excellent precision: $\hat{\alpha} = 4.6766$ (0.1983), $\hat{\beta} = 0.5620$ (0.0070), $\hat{\gamma} = 0.6079$ (0.0168), and $\hat{\delta} = -0.1771$ (0.0052).

Full Sample. Analogously to Figure 2, Figure 6 shows cross-sections of the surface estimated for the full sample in Figure 1 (right panel), used in forecasting energy consumption of the developing group countries. Aside from the poorest incomes to the left of the figure, there is a nearly monotonic increase in the coefficient over time. An increasing coefficient suggests a decrease in energy efficiency with no evidence of suppression as in the case of the surface for only the developed countries.




Figure 6: **Cross-Sections for the Full Sample.** Cross-sections of the surface $\beta(r, x)$ from fixing $r = \{1971, 1980, 1990, 2000, 2015\}$.




Figure 7: **Estimated Factor and Its Loading for the Full Sample.** Estimated using FPCA.




Figure 8: **Extending the Factor Loading for the Full Sample.** The leading factor is extended along the time dimension as a linear function of time

The eigenvalue ratio test of Ahn and Horenstein (2013) selects only one factor, which contributes 98.8% of the variation in the surface estimated for all countries. Figure 7 shows this factor and its loading. The factor loading looks quite a bit like the first factor loading for the developed group, increasing until about the turn of the century. In fact, the shape of the factor also resembles that of the first factor for the developed group, but with the critical difference that the factor for the developing group crosses zero at a much lower income and its value levels out at a higher magnitude than does that of the developed group. The positive factor and nearly monotonically increasing loading explain the near monotonicity of the cross-sections in Figure 6.

Much like the first factor loading for the developed group, that for the surface estimated using the full sample is nearly linear over 2007-2015, but declining (Figure 7, right panel). We extend it to 2035 using a linear function of time with slope estimated over that period in Figure 8. This recent decline may reverse an otherwise steadily increasing coefficient for some countries in the developing group.

3.3 Conditional Forecasting

Forecasts $h > 0$ periods ahead conditional on a GDP scenario $x_{i,T+h}$ are implemented using

$$y_{i,T+h}^* = \hat{\alpha}_i^* + (\hat{\gamma}_i^* + \hat{\beta}^*(T+h, x_{i,T+h}))x_{i,T+h} \quad (2)$$

where $\hat{\alpha}_i^*$, $\hat{\gamma}_i^*$, and $\hat{\beta}^*$ are draws from bootstrapped distributions of the respective parameter estimates, as described in the appendix. In this way, we evaluate only estimation uncertainty but not uncertainty in future error sequences and their serial correlation. Ignoring serial correlation is justified when conducting long-term forecasts such as ours.

We conduct forecasts for countries in both the developed and developing groups, looking specifically at the Germany (developed), Italy (developed), US (developed), Brazil (developing), China (developing), and India (developing). We condition on per capita real GDP to 2035 based on OECD real GDP forecasts³ and World Bank population forecasts.⁴

By using the surface estimated for all countries for Brazil, China, and India, we are allowing countries in the developing group to “graduate” from this group and consume more like the developed group once their per capita income is similar to that of the latter. Otherwise, if countries in the developing group must remain in that group, the coefficient on income may never level out or decrease with income.

Figures 9 and 10 show the input premises (left panels) and conditional forecasts of log per capita final energy consumption (right panels). For the most part, incomes (real GDPs) are projected to increase roughly as they have in the past across all countries. Not so for populations. While population growth in the US is projected to be steady, likely due to immigration, as is that for India, growth in China and Brazil is projected to slow noticeably.

³Downloaded from <https://data.oecd.org/gdp/gdp-long-term-forecast.htm> on February 12, 2018.

⁴Downloaded from <https://databank.worldbank.org/source/population-estimates-and-projections> on February 12, 2018.




Figure 9: **Input Premises and Conditional Forecasts: Germany, Italy, and the US.** Input premises of real GDP and population (left panels) and energy consumption forecasts (right panels) to 2035 for Germany (top panels), Italy (middle panels), and the US (bottom panels). Vertical line signifies the beginning of the forecast period, 2016.




Figure 10: **Input Premises and Conditional Forecasts: Brazil, China, and India.** Input premises of real GDP and population (left panels) and energy consumption forecasts (right panels) to 2035 for Brazil (top panels), China (middle panels), and India (bottom panels). Vertical line signifies the beginning of the forecast period, 2016.

At the extreme, the populations of Germany and Italy are projected to shrink, which is rather surprising given evident growth in recent years.

The results of our forecasts of per capita energy consumption conditional on these inputs are not surprising for the US and Germany: declines in per capita energy consumption are projected to continue at roughly the same pace. Also not surprisingly, but with the opposite sign, increases in per capita energy consumption for China and India are projected to continue, albeit at a slower rate than over the last decade.

The result for Italy is unexpected. Italy sharply reversed an increasing trend in consumption at around the turn of the century. This reversal roughly corresponds to a spike in population over that period. Population is projected to decline to 2035 – another reversal – which suggests an increasing energy consumption per capita.

We speculate that these patterns in Italy are due to immigration and an aging population. A decade of substantial increase started in roughly 2004, when ten Eastern European countries formally joined the European Union, easing restrictions on emigration to Italy. The decline starting in 2014 is likely due to both an aging population, as in many developed countries, and a tightening of restrictions on non-EU immigration following the influx of immigrants from Syria and other Middle Eastern and North African countries that has recently affected much of Europe and Southern Europe in particular.

The result for Brazil is also unexpected, in the sense that we project a reversal of the recent increase. However, the reversal makes sense in the context of our groups. As Brazil's population growth slows but GDP growth is maintained or accelerates, the consumption pattern of Brazil, a country in the developing group, is projected to decline similarly to countries in the developed group. Brazil nicely illustrates the advantage of our semiparametric panel approach, without which such a forecast would be difficult.

4 Conclusion

We introduce a nonparametrically specified, time-varying approach to forecasting regional energy consumption by way of a functional coefficient panel data model. As noted extensively in the extant literature, nonlinearity and time instability are key features of the relationship between income and energy consumption. Allowing for these features leads to effective forecasting of the latter conditional on the former.

By using a semiparametrically estimated panel model, we disentangle heterogeneous regional slopes from trending components in the relationship between income and energy consumption common to all countries or just to developed countries. As a group, developed countries have tended towards increasing energy efficiency (decreasing energy intensity) since the late 1990's, while developing countries continue to decrease energy efficiency (increase energy intensity). Estimating a common surface for all countries allows a developing country to switch from a decrease to an increase in energy efficiency as its per capita income increases, perhaps correlated with economic diversification and sectoral shifts to less energy-intensive industries.

Forecasts conditional on the input premises to 2035 from OECD and World Bank projections suggest continuing decreases in per capita energy consumption for Germany and the US and continuing increase for China and India. Forecasts of per capita energy consumption in Italy and Brazil show reversals: from decreasing to increasing in Italy and from increasing to decreasing in Brazil.

Although we focus on total per capita energy consumption at the national level in the present analysis, the methodology is quite general. The methodology is currently in use to produce official conditional forecasts of electricity consumption for South Korea. It would be straightforward to adapt to ISOs in North America, for example, using international or

state-level data to estimate the surface $\beta(r, x)$. We leave this for future research.

References

- Ahn, S.C. and A.R. Horenstein (2013). Eigenvalue ratio test for the number of factors, *Econometrica* 81, 1203-1227.
- Aue, A., D. Dubart Norinho, and S. Hörmann (2015). On the prediction of functional time series, *Journal of the American Statistical Association* 110, 378-392.
- Besse, P., H. Cardot, and D. Stephenson (2000). Autoregressive forecasting of some functional climatic variations, *Scandinavian Journal of Statistics* 27, 673-687.
- Bosq, D. (2000). Linear Processes in Function Spaces, in: *Lecture Notes in Statistics*, Vol. 149. Springer, New York.
- Bradford, T. (2018). *The Energy System: Technology, Economics, Markets, and Policy*. MIT Press, Cambridge.
- Chang, Y., M. Choi and J.Y. Park (2019a). A factor model for functional time series, mimeograph.
- Chang, Y., Y. Choi, C.S. Kim, J.I. Miller, and J.Y. Park (2016a). Disentangling temporal patterns in elasticities: A functional coefficient panel analysis of electricity demand, *Energy Economics* 60, 232-243.
- Chang, Y., Y. Choi, C.S. Kim, J.I. Miller, and J.Y. Park (2019b). Common factors and heterogeneity in economic relationships: A functional coefficient panel approach, mimeograph.
- Chang, Y., R.K. Kaufmann, C.S. Kim, J.I. Miller, J.Y. Park, and S. Park (2019c). Evaluating trends in time series of distributions: a spatial fingerprint of human effects on climate, *Journal of Econometrics*, forthcoming.
- Chang, Y., C.S. Kim, and J.Y. Park (2016b). Nonstationarity in time series of state densities, *Journal of Econometrics* 192, 152-167.
- Csereklyei Z., M.d.M. Rubio Varas, and D.I. Stern (2016). Energy and economic growth: The stylized facts, *The Energy Journal* 37, 223-255.
- Fan, J. and T. Huang (2005). Profile likelihood inferences on semiparametric varying-coefficient partially linear models, *Bernoulli* 11, 1031-1057.
- Galli, R. (1998). The relationship between energy intensity and income levels: Forecasting long term energy demand in Asian emerging countries, *The Energy Journal* 19, 85-105.

- Hörmann, S. and P. Kokoszka (2012). Functional Time Series, in: Handbook of Statistics: Time Series Analysis – Methods and Applications, pp. 157-186. Elsevier, Amsterdam.
- Hörmann, S., L. Kidziński, and M. Hallin (2015). Dynamic functional principal components, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 77, 319-348.
- Hyndman, R.J. and M.S. Ullah (2007). Robust forecasting of mortality and fertility rates: A functional data approach, *Computational Statistics and Data Analysis* 51, 4942-4956.
- Judson, R.A., R. Schmalensee, and T.M. Stoker (1999). Economic development and the structure of the demand for commercial energy, *The Energy Journal* 20, 29-57.
- Kaufmann, R.K. (2004). The mechanisms for autonomous energy efficiency increases: A cointegration analysis of the US energy/GDP ratio, *The Energy Journal* 25, 63-86.
- Luzzati, T. and M. Orsini (2009). Investigating the energy-environmental Kuznets curve, *Energy* 34, 291-300.
- Medlock, K.B. and R. Soligo (2001). Economic development and end-use energy demand, *The Energy Journal* 22, 77-105.
- Nguyen-Van, P. (2010). Energy consumption and income: A semiparametric panel data analysis, *Energy Economics* 32, 557-563.
- Richmond, A.K. and R.K. Kaufmann (2006a). Energy prices and turning points: The relationship between income and energy use/carbon emissions, *The Energy Journal* 27, 157-180.
- Richmond, A.K. and R.K. Kaufmann (2006b). Is there a turning point in the relationship between income and energy use and/or carbon emissions? *Ecological Economics* 56, 176-189.
- United States Energy Information Administration (2013). U.S. energy intensity projected to continue its steady decline through 2040, *Today in Energy*, March 1, 2013.
- Webster, M., S. Paltsev, and J. Reilly (2008). Autonomous efficiency improvement or income elasticity of energy demand: Does it matter? *Energy Economics* 30, 2785-2798.

Supplementary Online Material

Forecasting Regional Long-Run Energy Demand: A Functional Coefficient Panel Approach

Yoosoon Chang, Yongok Choi, Chang Sik Kim, J. Isaac Miller, Joon Y. Park

A Bootstrapping the Fitted Residuals

To implement the bootstrap procedure, we allow the fitted residuals from equation (1) to have both a common component and an idiosyncratic component. We use a low-dimensional VAR to model the common factors and an AR for the idiosyncratic component. For countries in the developed group, this procedure is identical to that employed in our earlier work (Chang *et al.*, 2019b). However, the procedure differs slightly for the group of all countries. For the sake of completeness, we describe the general procedure.

A.1 General Procedure

For each group – i.e., for the developed group or for the full sample – let N denote the number of countries for which we observe data over the whole time span, 1971 to 2015, so that $T = 45$. For these countries, bootstrapped residuals and response variables are obtained using the following steps.

1. Obtain the residuals $\hat{\varepsilon}_{it}$ from the model in equation (1) and demean $\tilde{\varepsilon}_{it} = \hat{\varepsilon}_{it} - \overline{\hat{\varepsilon}_i}$ across years. Because we employ a non-standard semiparametric estimation method, the residuals do not have a mean of zero.
2. Perform principal component analysis based on the covariance matrix of $\tilde{\varepsilon}_{\cdot t} = (\tilde{\varepsilon}_{1t}, \tilde{\varepsilon}_{2t}, \dots, \tilde{\varepsilon}_{Nt})'$. The number of factors r is chosen by both eigenvalue ratio and growth ratio tests of Ahn and Horenstein (2013). Decompose $\tilde{\varepsilon}_{\cdot t}$ as $\tilde{\varepsilon}_{\cdot t} = \hat{\Lambda}\hat{g}_t + \hat{\eta}_{\cdot t}$, where $\hat{g}_t$ is a $r \times 1$ vector of factors, $\hat{\Lambda} = (\hat{\lambda}_1, \dots, \hat{\lambda}_N)'$ with $\hat{\lambda}_i$ defined to be $r \times 1$ vector of factor loadings for variable $(\tilde{\varepsilon}_i)$, and $\hat{\eta}_{\cdot t} = (\hat{\eta}_{1t}, \dots, \hat{\eta}_{Nt})'$ is a vector of idiosyncratic components of $\tilde{\varepsilon}_{\cdot t}$.
3. To bootstrap the common component of residuals, model $\hat{g}_t$ as a VAR,

$$\hat{g}_t = B_1\hat{g}_{t-1} + \dots + B_k\hat{g}_{t-k} + \epsilon_t \quad (\text{A.1})$$

with the VAR order k determined by BIC. Fit the VAR model to get $\hat{B}_1, \dots, \hat{B}_k, \hat{\epsilon}_t$. Resample from $(\hat{\epsilon}_t - \overline{\hat{\epsilon}})$ to obtain bootstrap samples (ϵ_t^*) , from which bootstrap samples (g_t^*) are obtained using

$$g_t^* = \hat{B}_1g_{t-1}^* + \dots + \hat{B}_kg_{t-k}^* + \epsilon_t^*, \quad (\text{A.2})$$

where $g_t^* = \hat{g}_t$ for $t = 1, \dots, k$.

4. To bootstrap the idiosyncratic components of residuals, model $\hat{\eta}_{it}$ as an AR,

$$\hat{\eta}_{it} = \alpha_{i1}\hat{\eta}_{it-1} + \cdots + \alpha_{ip_i}\hat{\eta}_{it-p_i} + e_{it}. \quad (\text{A.3})$$

with the VAR order p_i determined by BIC. Fit the AR model to get $\hat{\alpha}_{i1}, \dots, \hat{\alpha}_{ip_i}, \hat{e}_{it}$. Resample $(\hat{e}_{it} - \overline{\hat{e}_i})$ to obtain bootstrap samples (e_{it}^*) ,⁵ from which bootstrap samples (η_{it}^*) are obtained using

$$\eta_{it}^* = \hat{\alpha}_1^i \eta_{i,t-1}^* + \cdots + \hat{\alpha}_k^i \eta_{i,t-p_i}^* + e_{it}^*, \quad (\text{A.4})$$

where $\eta_{it}^* = \hat{\eta}_{it}$ for $t = 1, \dots, p_i$.

5. From equations (A.2) and (A.4), the bootstrapped response variables are given by $y_{it}^* = \hat{y}_{it} + \hat{\lambda}_i' g_t^* + \eta_{it}^*$ respectively.
6. Parameters are re-estimated for each of 1,000 bootstrap replications of y_{it}^* holding x_{it} fixed and the distributions of these parameter estimates are used to construct standard errors and are fed into equation (2).

A.2 Group-Specific Data Irregularities

There are 37 countries with a data span from 1971 to 2015 ($N = 37, T = 45$) in the developed group. These are as follows: Australia, Austria, Bahrain, Belgium, Bermuda, Bulgaria, Canada, Chile, Cuba, Cyprus, Denmark, Finland, France, Germany, Greece, Hong Kong, Hungary, Iceland, Ireland, Israel, Italy, Japan, (Republic of) Korea, Luxembourg, Mexico, Netherlands, New Zealand, Norway, Panama, Poland, Portugal, Romania, Spain, Sweden, Switzerland, UK, and US. For these countries the procedure above is implemented with $r = 4$ factors chosen, $k = 1$ lags chosen for the common component, and $p_i \in \{1, \dots, 5\}$ lags chosen for each country.

Nine countries in the developed group have data deficiencies. Kuwait is missing data following the Gulf War, over 1992-1994, which we linearly interpolate. The former Soviet, Czechoslovakian, and Yugoslavian countries in the developed group, Belarus, the Czech Republic, Estonia, Latvia, Lithuania, Russian, Slovakia, and Slovenia, are missing data prior to 1996, which we ignore. We project the fitted residuals from these nine countries onto the common factor obtained from the residuals of the other 37 countries, and we use the residuals of these projections as the idiosyncratic component $\hat{\eta}_{it}$ in step 4 of the bootstrap procedure for these countries.

There are 21 countries with data deficiencies in the developing group. Cambodia, Tonga, and Samoa, and Yemen are missing data until 1980, 1981, 1982, and 1989. Macedonia and Uzbekistan are missing data until 1990, and Namibia is missing data until 1991. Armenia,

⁵For both the developed group and the full sample, we omit outliers exceeding $(-0.06345, 0.054107)$ following Chang *et al.* (2019).

Azerbaijan, Bosnia and Herzegovina, Georgia, Croatia, Kyrgyzstan, Kazakhstan, Moldova, Serbia, Tajikistan, Turkmenistan, and Ukraine are missing data until 1996. Niger is missing data only in 1979, which we interpolate.

This leaves 150 countries in the full sample spanning the full span ($N = 37 + 113 = 150$, see Chang *et al.*, 2019b, for the full list). We conduct the procedure as above, with $r = 4$ factors chosen, $k = 1$ lags chosen for the common component, and $p_i \in \{1, \dots, 5\}$ lags chosen for each country. Similarly to the developed group, we project the fitted residuals from the $21 + 9 = 30$ countries with missing data onto the common factors estimated for the other N countries and use the residuals of these projections as the idiosyncratic component for each country.