Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

Common Trends and Country Specific Heterogeneities in Long-Run World Energy Consumption

CAMA Working Paper 4/2024
January 2024

Yoosoon Chang
Indiana University
Centre for Applied Macroeconomic Analysis, ANU

Yongok Choi
Chung-Ang University

Chang Sik Kim
Sungkyunkwan University

J. Isaac Miller
University of Missouri

Joon Y. Park
Indiana University

| THE AUSTRALIAN NATIONAL UNIVERSITY

Abstract

We employ a semiparametric functional coefficient panel approach to allow an economic relationship of interest to have both country-specific heterogeneity and a common component that may be nonlinear in the covariate and may vary over time. Surfaces of the common component of coefficients and partial derivatives (elasticities) are estimated and then decomposed by functional principal components, and we introduce a bootstrap-based procedure for inference on the loadings of the functional principal components. Applying this approach to national energy-GDP elasticities, we find that elasticities are driven by common components that are distinct across two groups of countries yet have leading functional principal components that share similarities. The groups roughly correspond to OECD and non-OECD countries, but we utilize a novel methodology to regroup countries based on common energy consumption patterns to minimize root mean squared error within groups. The common component of the group containing more developed countries has an additional functional principal component that decreases the elasticity of the wealthiest countries in recent decades.

Keywords

JEL Classification

Address for correspondence:

(E) cama.admin@anu.edu.au

| THE AUSTRALIAN NATIONAL UNIVERSITY

Common Trends and Country Specific Heterogeneities
in Long-Run World Energy Consumption[*]

Yoosoon Chang[†]   Yongok Choi[‡]   Chang Sik Kim[§]   J. Isaac Miller[¶]

Joon Y. Park[†]

Abstract

We employ a semiparametric functional coefficient panel approach to allow an economic relationship of interest to have both country-specific heterogeneity and a common component that may be nonlinear in the covariate and may vary over time. Surfaces of the common component of coefficients and partial derivatives (elasticities) are estimated and then decomposed by functional principal components, and we introduce a bootstrap-based procedure for inference on the loadings of the functional principal components. Applying this approach to national energy-GDP elasticities, we find that elasticities are driven by common components that are distinct across two groups of countries yet have leading functional principal components that share similarities. The groups roughly correspond to OECD and non-OECD countries, but we utilize a novel methodology to regroup countries based on common energy consumption patterns to minimize root mean squared error within groups. The common component of the group containing more developed countries has an additional functional principal component that decreases the elasticity of the wealthiest countries in recent decades.

[†]Department of Economics, Indiana University and CAMA

[‡]School of Economics, Chung-Ang University

[§]Department of Economics, Sungkyunkwan University

[¶]Department of Economics, University of Missouri

1 Introduction

In time series or cross-sectional regressions, a relationship of interest is typically assumed to be homogeneous. Even most panel regressions assume homogeneity up to individual-specific intercepts that may be specified as stochastic (random effects) or nonstochastic (fixed effects). On the other hand, in random coefficient panel regressions the relationship may be completely heterogeneous across individuals and/or time. Such an approach effectively decomposes the relationship of interest into individually heterogeneous and/or temporally heterogeneous components. The latter components constitute a stochastic trend common to all individuals, and it may be informative about the temporal evolution of the relationship of interest.

Consider a panel of primary energy consumption by country. Energy consumption patterns are of great interest to scientists in a world with a burgeoning population augmented by economic growth per capita. The finiteness of our primary energy sources and the apparent fragility of human and ecological systems to pollutants from consuming energy have pushed the study of energy demand to the forefront of the agendas of many academics and policy makers. Understanding energy consumption underlies models of climate change used to price carbon, project damages, and ascertain the need for adaptation, for example.

Figure 1 presents an illustration of how the residuals from such a decomposition might look for a simple regression of energy consumption on its mean. The left panel of the figure shows substantial cross-country heterogeneity and no discernible common patterns in per capita energy consumption over time. Yet the right panel shows a stochastic trend common to non-OECD countries that is almost monotonically increasing and another one common to OECD countries that is decreasing since 2004. The figure makes clear the advantage of examining common components (right panel), because without them (left panel) the sample paths are resistant to ocular analysis.

A downside of tradition panel regressions is that estimation of a common temporal component precludes inference about its temporal dynamics, because each time effect is effectively treated as an independent draw. At the other extreme, a linear time trend is easy to analyze and predict but may vastly oversimplify temporal dynamics. Allowing for time-varying parameters could be considered an intermediate approach, in the sense that temporal dynamics are flexibly but meaningfully modeled, and there is certainly no shortage of studies that take such an approach in the energy literature: Park and Zhao (2010), Inglesi-Lotz (2011), Arisoy and Ozturk (2014), Chang *et al.* (2014, 2016a, 2021), Liddle *et al.* (2020), *inter alia*.[1]

Another downside of the traditional panel approach is its inflexible parametric specification. The

---

[1] A wide variety of intermediate approaches exist, including stochastically time-varying coefficients (Cooley and Prescott, 1976; Newbold and Bos, 1985, *inter alia*), deterministically time-varying coefficients (Lin and Teräsvirta, 1994; Park and Hahn, 1999; Bierens and Martins, 2010, *inter alia*), and common factor models (Sims and Sargent, 1977; Chamberlain and Rothschild, 1983; Bai and Ng, 2004, *inter alia*).
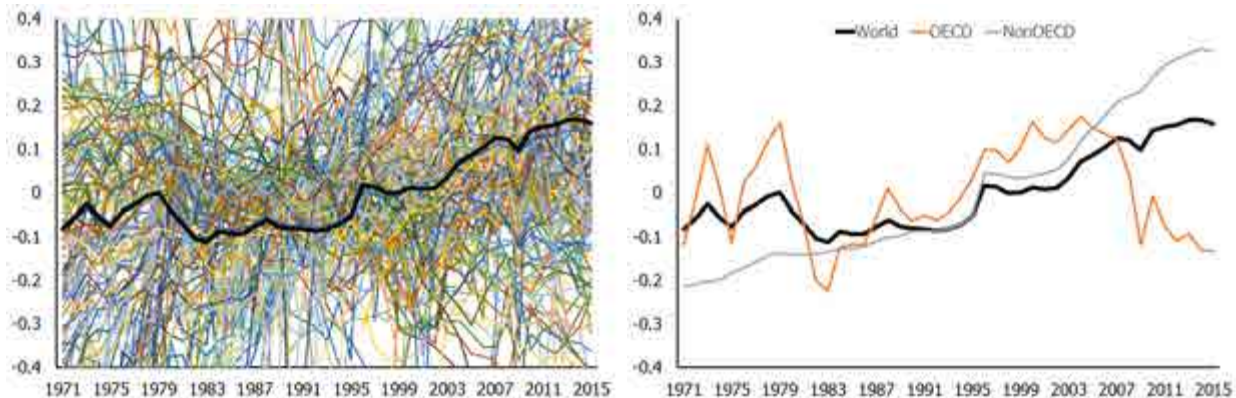
Figure 1: **Per Capita Energy Consumption.** Temporally demeaned energy consumption (millions of tonnes of oil equivalent, Mtoe) per capita over time by country (left panel) and aggregated for OECD countries, non-OECD countries, and the world (right panel). The large black line shows the world average in each panel.

empirical evidence for nonlinearity in the relationship between energy consumption and real GDP is well documented by Galli (1998), Judson *et al.* (1999), Medlock and Soligo (2001), Richmond and Kaufmann (2006), Chang *et al.* (2016a), Liddle and Huntington (2020), Liddle (2022), *inter alia.* However, many of these authors allow for nonlinearity by way of a quadratic function, which, although certainly an improvement over a linear specification, leaves much to be desired.

Our approach embraces both heterogeneity and nonlinearity of unknown form while explicitly modeling a common component to the coefficient and partial derivative (elasticity) – i.e., the economic relationship – of interest. We use fixed effects to account for individual heterogeneity in both slopes and intercepts, and we nonparametrically estimate a common component that varies smoothly over time and over the value of a regressor that may be correlated with time, specifically log real GDP, which proxies aggregate economic activity (output). In contrast with time series or cross-sectional settings – or even with simpler panel settings – a functional panel setting allows effective identification of a coefficient that is a function of two correlated arguments.

We further decompose the common components using functional principal component analysis into functional principal components (FPCs) or factors that are functions of economic activity and their loadings over time, and we employ a novel sieve-based bootstrap for interval estimates of loadings on the FPCs. Because we are extracting principal *components* from a common *components* of an economic relationship that is also a functional, we will henceforth refer to the latter either using the abbreviation FPC or as *factors* and reserve common *component* to describe the functional that is common to the economic relationship across individuals in the sample.

Although the research presented in this paper is firmly motivated by the need of economists and policy makers to better understand the long-run relationship between energy consumption and aggregate economic activity, the basic econometric issues raised and addressed by our modeling

approach are completely general. While the nonparametric estimator is related to that of extant studies, the decomposition of the relationship into idiosyncratic components and nonparametrically specified common components, which are further decomposed using functional principal component analysis, is novel both to the study of energy consumption and more generally, and it is the main contribution of this research.[2]

Turning back to energy consumption, much has been written on what some authors call the *energy elasticity*, *energy-GDP elasticity*, or *energy coefficient* (on aggregate production) of an economy, defined as the percentage increase in energy use per capita from a one percent increase in economic activity per capita (c.f. Brookes, 1972; Burke and Csereklyei, 2016, *inter alia*). Formally, the quantity may be expressed as a derivative dy/dx, where y is the log of energy consumption and x is the log of aggregate economic activity, typically measured using real GDP. A decrease in the energy coefficient is referred to as *autonomous energy efficiency improvement* (AEEI), which is a key parameter in most integrated assessment models used to model, project, and price the effects of anthropogenic climate change (Webster *et al.*, 2008), such as the MIT Emissions Prediction and Policy Analysis Model (Paltsev *et al.*, 2005; Chen *et al.*, 2022).

We find that a group of countries roughly corresponding to the OECD countries have elasticities with a common component of approximately 0.2 to 0.3 that has been declining in recent years. The decline is consistent with the pattern of energy intensity noted by Kaufmann (2004) and the US Energy Information Administration (2013) for the US, and by Webster *et al.* (2008) and Csereklyei *et al.* (2016) for developed countries more generally, and the range of the common component lies on the lower end of the spectrum of estimates of this elasticity in the extant literature. The remaining countries have a more complicated common pattern in their elasticities that generally appears to increase over time for fixed levels of economic activity over a range of roughly -0.2 to 0.4. Such patterns are certainly not obvious from Figure 1, and we believe such insights can only be gained by a panel approach as flexible as ours.

Further decomposing the common component of each group of countries into FPCs shows that both groups of countries have a leading FPC that is generally increasing over time and that we dub the *energy-increasing factor*. The group that contains mostly OECD countries has a secondary FPC that generally offsets the increase of the first FPC to the extent that the elasticity for the wealthiest countries generally has been declining during the latter part of the sample. We dub this factor the *intensity-reducing factor*.

Adding estimated idiosyncratic slopes on top of the common components reveals negative elasticities for some countries in both groups, a result which suggests the decoupling of economic growth and energy consumption for the most advanced economies in particular. Decoupling has been noted

---

[2]In a closely related paper (Chang *et al.*, 2021), we extend the loadings on the FPCs and bootstrap intervals in the time dimension for the purpose of out-of-sample forecasting.

by previous authors and is one of the "stylized facts" of Csereklyei *et al.* (2016), but our results suggest that decoupling may be more prevalent than previously thought.

The remainder of this paper is structured as follows. In Section 2, we explain the functional coefficient panel model, the semiparametric approach proposed to estimate the functional coefficient and elasticity surfaces, and the decomposition of these surfaces using functional principal component analysis. We discuss empirical estimates of the functional energy coefficient and elasticity surfaces, FPCs and loadings of the elasticity surfaces, and the elasticities of some specific countries in Section 3, and we conclude with Section 4. We describe the data, methodology for grouping the countries, and some additional details of our econometric procedures, including the bootstrap, in Appendices A, B, and C.

2 Model, Estimation, and Functional Principal Components

2.1 Model

We start by assuming that, for country i and year t with $i = 1, \ldots, N$ and $t = 1, \ldots, T$, log real energy consumption per capita Y_{it} is determined as

$$Y_{it} = A_{it} + B_{it} X_{it} \tag{1}$$

with log real GDP per capita X_{it} as a covariate, where A_{it} and B_{it} signify the intercept and slope coefficients, which are generally assumed to be stochastic and to vary non-deterministically over both i and t to allow for country- and time-specific components affecting energy consumption. Although our empirical analysis and our motivation for the modeling approach revolve around the relationship between energy consumption and economic activity (output) as measured by real GDP, this model is obviously and inherently completely general.

Of course, we need to further specify the intercept and slope coefficients A_{it} and B_{it} in equation (1) to enable empirical analysis. Denote by $\mathcal{F}_t$ the set of information on (X_{it}) accumulated up to time t across i. First, for the intercept A_{it}, we let

$$\mathbb{E}\left[A_{it} \middle| \mathcal{F}_t\right] = \alpha_i \tag{2}$$

for $i = 1, \ldots, N$ and $t = 1, \ldots, T$, which means that the part of energy consumption not related to real GDP is country-specific but does not vary over time. As in traditional panel models, (α_i) represent individual-specific fixed effects in our model. Specifically, each α_i measures the country-specific and time-invariant component of energy consumption.

Second, we specify the slope coefficient B_{it} as

$$\mathbb{E}\left[B_{it}\middle|\mathcal{F}_t\right] = \beta_i + \Gamma(t, X_{it}) \tag{3}$$

for $i = 1, \ldots, N$ and $t = 1, \ldots, T$. The specification in equation (3) distinguishes between country- and time-specific components of the slope coefficient. We define β_i to be the country-specific and time-invariant component of the slope for country i. The remainder of the slope coefficient is specified by $\Gamma(t, X_{it})$, a nonparametric function of time and real GDP. Time is included to account for generally trending behavior in the propensity to consume energy resulting from changes in energy consumption patterns, possibly due to sectoral shifts, changes in technology, and changes in preferences as the effects of anthropogenic climate change have become more evident. Real GDP is included as a proxy for the level of economic and social development of a country and to allow for nonlinearity in the relationship between Y_{it} and X_{it}. Here we set the slope coefficient of a country to be solely determined by its own current real GDP in addition to time. In particular, we assume that it is not affected by real GDP of a country's own past real GDP once its current real GDP is given. We believe that this is a reasonable modeling assumption, since we analyze the *long-run* energy demand for a country and abstract from the dynamics of its short-run energy demand.

Under our specifications in equations (2) and (3), equation (1) reduces to

$$Y_{it} = \alpha_i + [\beta_i + \Gamma(t, X_{it})]X_{it} + U_{it} \tag{4}$$

for $i = 1, \ldots, N$ and $t = 1, \ldots, T$, where U_{it} is an error term satisfying $\mathbb{E}[U_{it}|\mathcal{F}_t] = 0$. To allow for the general interactions of the arguments of the functional coefficient, we specify $\Gamma(t, X_{it})$ as an unknown function of t and X_{it}. Given our nonparametric specification of $\Gamma(t, X_{it})$, the country-specific components (β_i) are not identified for all $i = 1, \ldots, N$, so we impose the identifying restriction

$$\sum_{i=1}^{N} \beta_i = 0 \tag{5}$$

throughout the paper. The model in equation (4) with identifying restriction in equation (5) is semi-parametric and partially linear, having both a linear parametric part $\alpha_i + \beta_i X_{it}$ and a nonlinear and nonparametric part $\Gamma(t, X_{it})X_{it}$. It may therefore be viewed as a partially linear functional coefficient model with functional coefficient $\Gamma(t, X_{it})$.

Note that if we set $\Gamma(t, X_{it})$ equal to $\gamma_i X_{it}$, ignoring the time-varying trend in Γ and making the functional form of Γ linear, then our model simplifies to $Y_{it} = \alpha_i + \beta_i X_{it} + \gamma_i X_{it}^2 + U_{it}$. This quadratic specification is often fitted in the energy demand literature noted above, and, by swapping out energy consumption for its resulting emissions, the quadratic model becomes that often used to estimate the environmental Kuznets curve (EKC) in a related literature.

The model in equations (4) and (5) and may be viewed as a conditional expectation function model by writing

$$\mathbb{E}[Y_{it}|t = r, X_{it} = x] = \alpha_i + \beta_i x + \Gamma(r, x)x.$$

to achieve $\mathbb{E}[U_{it}|\mathcal{F}_t] = 0$. Therefore, if we set $\Phi_i(r, x) = (\partial/\partial x)\mathbb{E}\left[Y_{it}|t = r, X_{it} = x\right]$, it follows that

$$\Phi_i(r, x) = \beta_i + \Phi(r, x) = \beta_i + \Gamma(r, x) + \Gamma_x(r, x)x \tag{6}$$

where Γ_x denotes the partial derivative of Γ with respect to x. $\Phi_i(r, x)$ clearly represents the energy-GDP elasticity introduced above and discussed in more detail below, at time r and at log real GDP x for country $i = 1, \ldots, N$, and $\Phi(r, x)$ represents its common component. In light of the identifying restriction in equation (5), country-specific components of the elasticity are zero on average, so that the surface $\Phi(r, x) = \Gamma(r, x) + \Gamma_x(r, x)x$ may be interpreted as an average elasticity at the point (r, x). However, the term "average" applies very loosely here, because it may very well be that no country has a log real GDP of x at time r.

2.2 Estimation

We employ a two-step estimation procedure to estimate $\Gamma(r, x)$ at any point (r, x) adapted from that of Fan and Huang (2005) and related to that of Cai (2007), Cai *et al.* (2009), and Chang *et al.* (2016a, 2021). We introduce the weight function

$$W_{it}(r, x) = \frac{1}{h_r h_x}K\left(\frac{t - r}{h_r}\right)K\left(\frac{X_{it} - x}{h_x}\right)\varpi_{it}, \tag{7}$$

where K is a kernel function, h_r and h_x are bandwidth parameters for r and x respectively, and $\varpi_{it} > 0$ is such that $\sum_{i=1}^{N}\varpi_{it} = 1$ for each t. The weight function $W_{it}(r, x)$ is identical to that of Chang *et al.* (2016a) up to multiplication by ϖ_{it}, which we discuss in Section 3. We utilize a standard normal kernel function, but we do not expect the choice of kernel function to make any substantive difference. Bandwidth selection is described in more detail in Appendix B.1.

Suppose for the moment that α_i and β_i are known for all i. In that case, we could estimate $\Gamma(r, x)$ for any point (r, x) using the kernel regression

$$\tilde{\Gamma}(r, x) = \left[\sum_{i=1}^{N}\sum_{t=1}^{T}X_{it}^2 W_{it}(r, x)\right]^{-1}\left[\sum_{i=1}^{N}\sum_{t=1}^{T}X_{it}(Y_{it} - \alpha_i - \beta_i X_{it})W_{it}(r, x)\right]$$
$$= \hat{\Gamma}_0(r, x) - \sum_{i=1}^{N}\alpha_i\hat{\Gamma}_{1i}(r, x) - \sum_{i=1}^{N}\beta_i\hat{\Gamma}_{2i}(r, x),$$

where

$$\hat{\Gamma}_0(r,x) = \left[\sum_{i=1}^{N}\sum_{t=1}^{T} X_{it}^2 W_{it}(r,x)\right]^{-1}\left[\sum_{i=1}^{N}\sum_{t=1}^{T} X_{it}Y_{it}W_{it}(r,x)\right]$$

$$\hat{\Gamma}_{1i}(r,x) = \left[\sum_{i=1}^{N}\sum_{t=1}^{T} X_{it}^2 W_{it}(r,x)\right]^{-1}\left[\sum_{t=1}^{T} X_{it}W_{it}(r,x)\right]$$

$$\hat{\Gamma}_{2i}(r,x) = \left[\sum_{i=1}^{N}\sum_{t=1}^{T} X_{it}^2 W_{it}(r,x)\right]^{-1}\left[\sum_{t=1}^{T} X_{it}^2 W_{it}(r,x)\right].$$

Note that $\hat{\Gamma}_0(r,x)$ is the nonparametric kernel estimator from the regression of (Y_{it}) onto (X_{it}) for $i = 1, \ldots, N$ and $t = 1, \ldots, T$ and, for each i, $\hat{\Gamma}_{1i}(r,x)$ and $\hat{\Gamma}_{2i}(r,x)$ are the nonparametric kernel estimators from regressions on (X_{it}) of an indicator for country i and of the product of that indicator with (X_{it}), respectively.

Subtract $\tilde{\Gamma}(t, X_{it})$ from both sides of the model in equation (4) and group terms to obtain

$$\left[Y_{it} - \hat{\Gamma}_0(t, X_{it})X_{it}\right] = \sum_{i=1}^{N} \alpha_i \left[1 - \hat{\Gamma}_{1i}(t, X_{it})\right] + \sum_{i=1}^{N} \beta_i \left[1 - \hat{\Gamma}_{2i}(t, X_{it})\right] X_{it} + \tilde{U}_{it}, \qquad (8)$$

where $\tilde{U}_{it} = U_{it} - [\tilde{\Gamma}(t, X_{it}) - \Gamma(t, X_{it})]X_{it}$. Given the feasible estimators $\hat{\Gamma}_0(t, X_{it})$, $\hat{\Gamma}_{1i}(t, X_{it})$, and $\hat{\Gamma}_{2i}(t, X_{it})$, the parameters α_i and β_i are estimated from equation (8) using N least squares regressions for each $i = 1, \ldots, N$. These regressions comprise the first step of the two-step estimation procedure.

Once the least squares estimators $\hat{\alpha}_i$ and $\hat{\beta}_i$ are obtained in the first step, the second step is simply to calculate the common functional coefficient using

$$\hat{\Gamma}(r,x) = \hat{\Gamma}_0(r,x) - \sum_{i=1}^{N} \hat{\alpha}_i \hat{\Gamma}_{1i}(r,x) - \sum_{i=1}^{N} \hat{\beta}_i \hat{\Gamma}_{2i}(r,x),$$

for any point (r,x), which is a feasible version of the infeasible estimator $\tilde{\Gamma}(r,x)$. Following reasoning similar to that of Chang *et al.* (2016a) for a related partially linear panel model, we may expect the estimators $\hat{\alpha}_i$, $\hat{\beta}_i$, and $\hat{\Gamma}(r,x)$ to be consistent under general conditions as $T \to \infty$ for fixed N.

If an estimate of $\Gamma_x(r,x)$ is obtained, the energy-GDP elasticity $\Phi_i(r,x)$ may be estimated for any country $i = 1, \ldots, N$ and at any point (r,x) from equation (6) using estimates $\hat{\beta}_i$, $\hat{\Gamma}(r,x)$, and $\hat{\Gamma}_x(r,x)$. For the estimator of $\Gamma_x(r,x)$, we employ

$$\hat{\Gamma}_x(r,x) = \frac{1}{2\delta}(\hat{\Gamma}(r, x + \delta) - \hat{\Gamma}(r, x - \delta)),$$

using some small increment δ.[3] Clearly, an alternative estimator defined as

$$\hat{\Gamma}_x(r,x) = \frac{\partial}{\partial x}\hat{\Gamma}_0(r,x) - \sum_{i=1}^{N}\hat{\alpha}_i\frac{\partial}{\partial x}\hat{\Gamma}_{1i}(r,x) - \sum_{i=1}^{N}\hat{\beta}_i\frac{\partial}{\partial x}\hat{\Gamma}_{2i}(r,x),$$

which we obtain by taking analytical derivatives of $\hat{\Gamma}_0(r,x), \hat{\Gamma}_{1i}(r,x)$ and $\hat{\Gamma}_{2i}(r,x)$, may also be used.[4]

2.3 Functional Factor Analysis

To analyze the variation over time of the common component $\Gamma(r,x)$ of the coefficient on log real GDP or the common component $\Phi(r,x)$ of the energy-GDP elasticity, we regard $\Gamma(t,\cdot)$ or $\Phi(t,\cdot)$ as a functional time series – i.e., a time series of functions of log real GDP, whose observations are available for $t = 1,\dots,T$. This functional approach is very useful to investigate how the coefficient or elasticity changes over time as functions of log real GDP. Our empirical study uses the functional factor model of Chang *et al.* (2021, 2023), which allows us to represent and interpret a functional time series effectively as a linear combination of a small number of functional factors.

For our functional factor analysis, we let (f_t) be a functional time series defined as

$$f_t = \Gamma(t,\cdot) - \overline{\Gamma}(\cdot) \quad \text{or} \quad \Phi(t,\cdot) - \overline{\Phi}(\cdot)$$

for $t = 1,\dots,T$, where $\overline{\Gamma}(\cdot) = T^{-1}\sum_{t=1}^{T}\Gamma(t,\cdot)$ and $\overline{\Phi}(\cdot) = T^{-1}\sum_{t=1}^{T}\Phi(t,\cdot)$, and regard them formally as random elements taking values in a Hilbert space H of square integrable functions on $\mathbb{R}$ endowed with inner product $\langle u,v\rangle = \int u(x)v(x)dx$ for $u,v \in H$ and norm $\|v\|^2 = \langle v,v\rangle = \int v^2(x)dx$. Then we consider the minimization problem

$$\sum_{t=1}^{T}\left\|f_t - \sum_{k=1}^{K}c_{kt}\phi_k\right\|^2,$$

with respect to a set of constants (c_{kt}) in $\mathbb{R}$ for $k = 1,\dots,K$ and $t = 1,\dots,T$ and an orthonormal set of square integrable functions (ϕ_k) in H for $k = 1,\dots,K$.

The solution to this minimization problem is given by $(\hat{c}_{kt})$ and $(\hat{\phi}_k)$ for (c_{kt}) and (ϕ_k), respectively, where $(\hat{\phi}_k)$ are the K-leading functional principal components (FPCs) of (f_t) and $(\hat{c}_{kt})$

---

[3]We use increments δ on supports $[\ln(2000),\ln(80000)]$ for the group that includes most OECD countries, which we subsequently label intensity-reducing, and $[\ln(500),\ln(20000)]$ for the group that includes most non-OECD countries in our sample, which we subsequently label intensity-increasing, where δ is equal to the difference of the log of increments of 100 real USD and thus varies over the level of real GDP.

[4]This alternative estimator, however, doesn't seem to be as reliable as the one we obtain from the numerical derivative defined above. Therefore, we use the numerical derivative in our empirical analysis in the paper.

are their loadings given by $\hat{c}_{kt} = \langle \hat{\phi}_k, f_t \rangle$, which will be defined more precisely in our subsequent discussions. It follows that

$$f_t \approx \sum_{k=1}^{K} \hat{c}_{kt} \hat{\phi}_k$$

provides the best approximation of (f_t) as a linear combination of K-dimensional basis functions. This approximation may indeed be obtained directly by the projection of (f_t) on the K-dimensional subspace $\hat{H}_K$ of H, which is spanned by the orthonormal basis $(\hat{\phi}_k)$ consisting of K-leading FPCs.

If we let

$$f_t = \sum_{k=1}^{K} c_{kt} \phi_k + e_t \tag{9}$$

with $c_t = (c_{1t}, \ldots, c_{Kt})' \in R_K$ for $t = 1, \ldots, T$ and an orthonormal basis (ϕ_k) for $k = 1, \ldots, K$ spanning H_K, and assume that[5]

$$\frac{1}{T} \sum_{t=1}^{T} c_t c_t' \to_p Q > 0 \quad \text{and} \quad \frac{1}{T} \sum_{t=1}^{T} (e_t \otimes e_t) \to_p 0, \tag{10}$$

we may *formally* define (ϕ_k) as functional factors and (c_{kt}) as their loadings in the sense of Chang *et al.* (2023). Our specification in (9) and (10) implies that (ϕ_k) appear in (f_t) pervasively over time, while (e_t) contributes to (f_t) only sporadically over time, both of which are given as functions of log real GDP. Therefore, we may identify the functional components $\left(\sum_{k=1}^{K} c_{kt} \phi_k \right)$ and (e_t) in (9) as the common and idiosyncratic components of (f_t), respectively, and define (ϕ_k) as functional factors. This is completely analogous to the identification and definition of factors in conventional scalar factor models.

Under our specification, the common component $\left(\sum_{k=1}^{K} c_{kt} \phi_k \right)$ thus fully represents all regular patterns in the fluctuations of (f_t) over time. The first condition in (10) implies that each functional factor ϕ_k representing a common functional feature in (f_t) appears regularly over time for $k = 1, \ldots, K$.[6] The second condition in (10), on the other hand, requires that the idiosyncratic component (e_t), which is introduced to allow for the presence of various temporary and time-specific features in (f_t), affects (f_t) only sporadically over time.[7] The functional factors (ϕ_k) and their loadings (c_{kt}) are consistently estimated by $(\hat{\phi}_k)$ and $(\hat{c}_{kt})$, respectively, under mild technical

---

[5]Here and in some of our subsequent discussions, we use the tensor product $u \otimes v$ of functions u and v in H, which is an operator (of rank 1) defined as $(u \otimes v)w = \langle v, w \rangle u$ for all $w \in H$. Intuitively, we may view the tensor product of functions in H as a generalization of the outer product of vectors in $\mathbb{R}^n$. In fact, if u and v are vectors in $\mathbb{R}^n$, $u \otimes v$ reduces to uv'.

[6]This implies, in particular, that $T^{-1} \sum_{t=1}^{T} \langle \phi_k, f_t \rangle^2 \to_p \sigma_k^2 > 0$ for all $k = 1, \ldots, K$.

[7]This implies that (e_t) is not a regular stationary time series, for which we would expect $T^{-1} \sum_{t=1}^{T} (e_t \otimes e_t) \to_p \Omega$, where $\Omega = \mathbb{E}(e_t \otimes e_t)$. In our functional factor model, (e_t) is defined to be the idiosyncratic component of (f_t) in the sense that it is the component of (f_t) which appears not regularly but only intermittently over time. As a result, the variation in (e_t) is not accumulated as fast as T, which yields $T^{-1} \sum_{t=1}^{T} (e_t \otimes e_t) \to_p 0$, as specified in (10).

conditions. Moreover, the number K of functional factors can be found using the extension by Chang *et al.* (2023) of the eigenvalue ratio test of Ahn and Horenstein (2013) for the standard factor model to the functional factor model used in our paper. The reader is referred to Chang *et al.* (2023) for more details. The reader interested in more details of standard functional data analysis is referred to Ramsay and Silverman (2005) and the references therein and especially Bosq (2000) for an introduction to functional time series.

For any given $K \leqslant T$, the K-leading FPCs $(\hat{\phi}_k)_{k=1}^{K}$ of $(f_t)_{t=1}^{T}$ can be readily obtained as the eigenfuctions of the sample variance operator

$$\hat{\Sigma} = \frac{1}{T} \sum_{t=1}^{T} (f_t \otimes f_t) \tag{11}$$

of (f_t), associated with its K-largest eigenvalues $\lambda_1 > \cdots > \lambda_K$.[8] Once the FPCs $(\hat{\phi}_k)_{k=1}^{K}$ are found, their loadings $(\langle \hat{\phi}_k, f_t \rangle)$ may easily be obtained. For more details of functional principal component analysis, the reader is referred to Ramsay and Silverman (2005) among others.[9]

As discussed, the FPCs $(\hat{\phi}_k)$ are functions of log real GDP. The estimated leading FPC $\hat{\phi}_1$ identifies the function of log real GDP that is associated with the largest proportion of fluctuations of (f_t) over time. The actual proportion of temporal fluctuations of (f_t) explained by $\hat{\phi}_1$ is given by

$$\pi_1 = \frac{\lambda_1}{\sum_{k=1}^{T} \lambda_k}. \tag{12}$$

The secondary FPC $\hat{\phi}_2$ represents the function of log real GDP that is orthogonal to $\hat{\phi}_1$ and explains the second largest proportion of fluctuations of (f_t) over time. The proportion π_2 of temporal fluctuations of (f_t) explained by $\hat{\phi}_2$ can be obtained similarly to π_1 in (12) using the ratio of λ_2 and $\sum_{k=1}^{T} \lambda_k$. In particular, the total proportion of temporal fluctuations of (f_t) accounted for by the K-leading FPCs is given by $\sum_{k=1}^{K} \pi_k$. Our functional factor model implies that this much of the variation in (f_t) is driven by the influences of some functional factors repeated over time, while the rest of the variation in (f_t) is originated from other sources that are of temporary and intermittent nature.

It turns out that the factor structures of (f_t) are low-dimensional for energy coefficients and elasticities as will be discussed in the next section. The eigenvalue ratio test finds only one or

---

[8]For the actual computation of (ϕ_k), we do not need to use the sample variance operator of (f_t) defined in (11), which is abstract and difficult to deal with. We may obtain the FPCs $(\hat{\phi}_k)$ directly from the eigenvectors and eigenvalues of the Gram matrix of (f_t), which is a $T \times T$ matrix. See Chang *et al.* (2023) for more details.

[9]The terminology used in the functional data analysis literature is not consistent. In Ramsay and Silverman (2005), our FPCs $(\hat{\phi}_k)$ and their loadings $(\langle \hat{\phi}_k, f_t \rangle)$ are referred to as principal component functions and principal component scores of functional data (f_t), respectively. In Wikipedia, on the other hand, the loadings $(\langle \hat{\phi}_k, f_t \rangle)$ are defined as the principal components of functional data (f_t).

two functional factors, and they explain more than 80% of the temporal variations. This factor structure greatly simplifies the analysis of historical energy consumption and makes it possible to effectively forecast the future energy demand, as Chang *et al.* (2021) do.

3 Energy-GDP Elasticity Estimation and Decomposition

For energy consumption z at GDP level w, energy-GDP elasticity dy/dx with $y = \ln z$ and $x = \ln w$ is closely related to *energy intensity z/w* and *energy efficiency dz/dw*, which are expressed in units of energy consumed per dollar instead of a unitless elasticity. A negative value is sometimes referred to as an *autonomous energy efficiency increase* (AEEI) (Kaufmann, 2004), where autonomous emphasizes that the change is independent of covariates, such as price. AEEI is an important component of integrated assessment models such as the MIT Emissions Prediction and Policy Analysis Model (Paltsev *et al.*, 2005; Chen *et al.*, 2022) used by economists to study climate change.

The terminology "elasticity" seems ubiquitous in this literature but is somewhat problematic. Adams and Miovic (1968) make a distinction between "gross [energy] elasticity," which is measured in this way, and "'true' energy elasticity," which is a parameter of a structural model. They note that the former underestimates the latter. The distinction between these types of elasticities is evident from the fact that regression techniques are often used to estimate a derivative. Without energy prices and other structural demand components, regression-based techniques suffer from an omitted variable bias in estimating a structural elasticity, a partial rather than total derivative. Adams and Miovic (1968) emphasize the need to examine relative efficiency of energy inputs and Chang *et al.* (2016a) emphasize the difference in interpretation of the elasticity across economic sectors when aggregate consumption data are used.

In addition to "elasticity," Brookes (1972) uses "coefficient" for $dy/dx = d\ln z / d\ln w$, implicitly reflecting both the regression approach to estimation and the direction of causality later explored empirically by Kraft and Kraft (1978) *inter alia*. However, this terminology is also problematic, because Adams and Miovic (1968) and even Brookes (1972) note instability in this derivative as the composition of a country's energy consumption bundle changes over time. Galli (1998), Judson *et al.* (1999), Medlock and Soligo (2001), Richmond and Kaufmann (2006), *inter alia* find nonlinear relationships between these variables, which undermines the interpretation of the coefficient itself as a derivative.

We favor the terminology *energy-GDP elasticity* (like "gross [energy] elasticity" of Adams and Miovic, 1968) recently employed by Burke and Csereklyei (2016) for the derivative $dy/dx = d\ln z / d\ln w$, and estimation of this (total) derivative is the main aim of our empirical analysis. We reserve *coefficient* for the coefficient on $x = \ln w$, which is itself a function of x.

Our econometric model allows the energy-GDP elasticity to vary over time and over different levels of GDP, and we apply it to a sample of countries that covers over 90% of the global population since 1971. In our estimation procedure, ϖ_{it} in the definition of $W_{it}(r,x)$ in equation (7) is the proportion of country i's population to the global population in year t, and it is introduced to account for difference in population sizes across countries.

A novelty of our econometric approach is that we allow the elasticity to follow a common path in each of the two subsets of countries while countries retain idiosyncratic components that reflect historical, political, cultural, geographical, or geological differences affecting how energy is consumed. Countries are grouped by commonality in energy intensity patterns over time. At the risk of oversimplification, we label one group (mostly OECD countries) the *energy intensity-reducing* (or simply *intensity-reducing*) group and the other group (mostly non-OECD countries) the *energy intensity-increasing* (or simply *intensity-increasing*) group. Indeed, our grouping procedure, which is described in detail in Appendix A.2, starts with the OECD vs. non-OECD distinction and strategically switches countries between groups to minimize root mean squared error, but the final grouping is based on commonalities in the functional coefficient on log real GDP rather than a political grouping. As will become clear shortly, elasticities of the intensity-increasing group are driven by what we call the intensity-increasing factor, while those of the intensity-declining group are driven by a similar intensity-increasing factor tempered by what we call the intensity-reducing factor.

The labels above describe commonalities given by $\Gamma(r,x)$ in how the energy intensity coefficients for each group have changed over time and especially in the past decade or so, but they do not necessarily describe the coefficients themselves or the energy-GDP elasticities of individual countries, both of which include idiosyncratic components (β_i). An intensity-reducing country such as Iceland or South Korea may have a high energy intensity from large positive β_i. In the case of Iceland, the high intensity results from abundant hydroelectric and geothermal energy sources rather than technical inefficiency, as Bradford (2018) points out. South Korea is one of the pioneers of energy-intensive industrial export-driven growth. On the other hand, an intensity-increasing country with a developed non-energy-intensive service sector such as South Africa may have a low energy intensity from a large negative β_i.

3.1 Functional Coefficient Estimates

Figure 2 shows the estimated functional coefficient surface $\hat{\Gamma}(r,x)$ for each group. The nonparametric procedure allows us to estimate the surfaces for any value of r and x using nearby observations of t and X_{it}. However, the precision of the estimated surface at a given pair of r and x is expected to depend on the actual frequency of X_{it} in its vicinity. If observations are densely (sparsely) populated in the neighborhood of the point given by a pair of r and x, we may expect the estimated
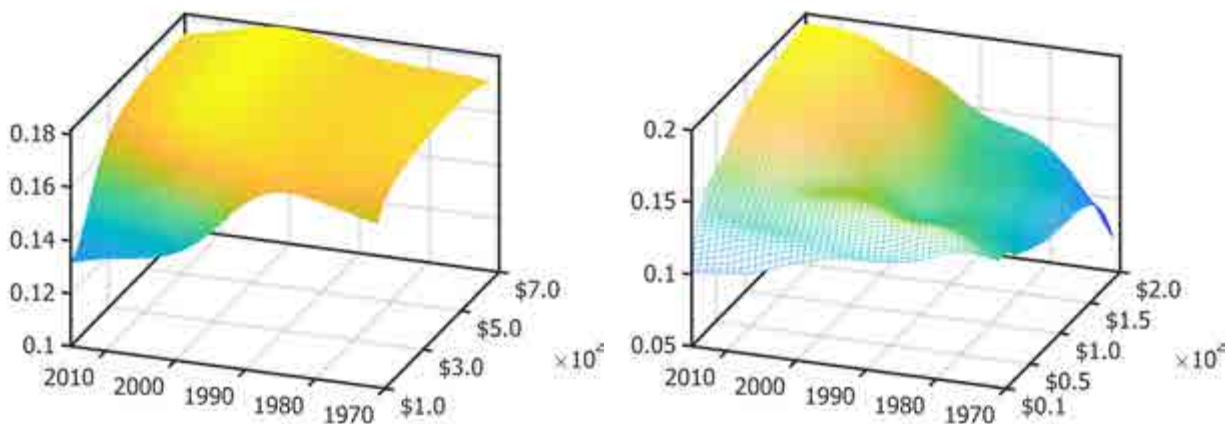
Figure 2: **Estimated Functional Coefficient Surfaces.** Estimates of the common components $\Gamma(r, x)$ of the coefficients on log real GDP for the energy intensity-reducing group (left) and the energy intensity-increasing group (right). The surfaces are bias-corrected using the bootstrap procedure described in Appendix C.

surface at that point to be more (less) precise. Dark blue areas in the lower left and upper right corners of the surfaces in the figure are likely estimated with less precision than yellow and orange areas for this reason. There simply were not very many countries with output levels near $5,000$ per capita in the first (wealthier) group towards the end of the sample and with output levels near $20,000$ per capita in the second (poorer) group towards the beginning of the sample.

At higher real GDP levels – approximately more than $30,000 per capita and especially noticeable at the maximum of $70,000 per capita – a common pattern is apparent for the energy intensity-reducing group. The functional coefficient clearly increases with time up to about the turn of the century and then decreases. The figure also suggests an increase in the coefficient with real GDP for any fixed year, which provides some justification for a quadratic approximation used by previous authors. Keeping in mind that countries' real GDPs tend to grow over time, a country will tend to move from the lower right to the upper left of the surface. In this light, a coefficient that decreases as a country's economy grows and hence generates an inverted "U" shape is possible but can result only from conflating economic growth with the natural passage of time.

Looking at the estimated functional coefficient surface of the energy intensity-increasing group in the right panel of Figure 2 and keeping in mind the caveat that we should discount the lower left and upper right corners of the surfaces, we see that the coefficient is mostly increasing in both time and real GDP. Moving from lower right to upper left gives the impression of an unambiguous increase in the functional coefficient for countries in this group. Indeed, countries in this group that have developed rapidly, like China, have generally done so by relying on energy-intensive industrial sectors as the engines for their growth. It does not appear to matter whether economic growth and

	α_i	s.e.		β_i	s.e.	
		asym.	bstp.		asym.	bstp.
Iceland	-6.75	1.02	2.04	0.65	0.07	0.20
Israel	-7.38	0.70	0.65	0.60	0.04	0.06
South Korea	-6.72	0.48	0.41	0.60	0.04	0.04
Portugal	-7.00	0.66	0.94	0.58	0.05	0.10
Chile	-6.36	0.74	0.95	0.54	0.07	0.10
Greece	-6.33	0.64	0.37	0.51	0.04	0.04
Mexico	-5.75	1.37	1.24	0.46	0.13	0.12
New Zealand	-3.85	0.82	0.42	0.31	0.05	0.03
Hong Kong	-4.68	0.47	0.66	0.28	0.03	0.06
Spain	-3.78	0.62	0.78	0.26	0.04	0.08
$\vdots$				$\vdots$		
Russia	1.84	0.93	1.33	-0.23	0.09	0.13
Switzerland	1.97	0.82	0.44	-0.27	0.05	0.04
Canada	2.77	0.85	0.53	-0.27	0.05	0.04
France	1.96	0.85	0.49	-0.27	0.05	0.03
Germany	2.98	0.67	0.53	-0.36	0.03	0.05
Denmark	3.02	0.72	0.39	-0.37	0.04	0.03
United Kingdom	3.39	0.72	0.49	-0.42	0.04	0.04
Sweden	4.06	0.67	0.49	-0.44	0.03	0.03
United States	5.27	1.17	0.86	-0.52	0.08	0.07
Cuba	3.47	0.95	1.50	-0.54	0.09	0.15

Table 1: **Selected Estimated Heterogeneous Components: Energy Intensity-Reducing Group.** Estimates of α_i and β_i for countries with ten highest and lowest estimates of β_i. Asymptotic standard errors denoted by "asym." and bootstrapped standard errors denoted by "bstp." and calculated as described in Appendix C.

time are conflated: we do not expect an inverted "U" shape in the coefficient on economic activity for any of the countries in this group.

Table 1 shows a subset of the results for the individual intercept and slope coefficients $\hat{\alpha}_i$ and $\hat{\beta}_i$ estimated for the energy intensity-reducing group. Specifically, the countries with the ten largest and ten smallest estimated values of β_i are shown. We find a strong negative correlation between $(\hat{\alpha}_i)$ and $(\hat{\beta}_i)$, and the countries with large values of $(\hat{\beta}_i)$ tend to have small values of $(\hat{\alpha}_i)$. This negative correlation is analyzed in Appendix B.2. In interpreting the results, one should keep in mind the identifying restriction in equation (5) that the idiosyncratic slopes sum to zero. To an approximation, we can interpret a negative $\hat{\beta}_i$ to mean that the energy intensity of a country is declining more rapidly or growing less rapidly than the intensities of the rest of the intensity-

	α_i	s.e.		β_i	s.e.	
		asym.	bstp.		asym.	bstp.
Afghanistan	-16.15	1.68	1.84	1.70	0.24	0.26
Haiti	-7.97	1.71	1.32	0.75	0.23	0.18
Thailand	-6.41	0.46	0.46	0.52	0.03	0.05
Malaysia	-6.10	0.52	0.45	0.48	0.05	0.05
Taiwan	-5.64	0.77	0.35	0.47	0.07	0.04
Bolivia	-5.79	0.51	0.74	0.44	0.05	0.09
Algeria	-6.28	1.55	1.43	0.44	0.16	0.16
Morocco	-5.21	0.48	0.63	0.32	0.04	0.07
Vietnam	-4.47	0.31	0.23	0.30	0.03	0.02
Malawi	-5.44	1.61	1.13	0.30	0.23	0.16
$\vdots$				$\vdots$		
Cameroon	1.03	0.68	0.81	-0.42	0.08	0.10
Rwanda	0.77	1.61	0.85	-0.50	0.23	0.12
Chad	0.56	0.87	1.13	-0.52	0.12	0.15
Burkina Faso	1.11	0.46	0.67	-0.54	0.06	0.10
Colombia	4.07	0.66	1.14	-0.69	0.06	0.13
Philippines	3.40	0.82	0.76	-0.71	0.09	0.09
South Africa	5.51	1.61	3.09	-0.73	0.17	0.33
Niger	2.23	0.72	0.84	-0.77	0.10	0.12
Uzbekistan	5.92	0.57	0.96	-0.83	0.06	0.11
Burundi	3.76	1.03	1.21	-0.95	0.16	0.18

Table 2: **Selected Estimated Heterogeneous Components: Energy Intensity-Increasing Group.** Estimates of α_i and β_i for countries with ten highest and lowest estimates of β_i. Asymptotic standard errors denoted by "asym." and bootstrapped standard errors denoted by "bstp." and calculated as described in Appendix C.

reducing group, while we can interpret a positive $\hat{\beta}_i$ to mean that energy intensity is declining less rapidly or increasing more rapidly than the intensities of rest of the group.

Iceland and South Korea rank among the most intensive users in the intensity-reducing group. As noted above, Iceland's intensity likely results from abundant hydroelectric and geothermal energy sources. Other large producers of geothermal energy on this list include Portugal and Greece, while Mexico and New Zealand have large installed geothermal capacities. Also noted above, South Korea's intensity likely results from its energy-intensive export-driven economy, as does the intensity of Hong Kong, which is treated separately from China in the data because it was separate over most of the sample. Israel's presence at this end of the spectrum is more difficult to explain, but we speculate that its geopolitical isolation make energy efficiency paramount, hence the low $\hat{\alpha}_i$, yet

its development of an energy-intensive defense industry dictates a less negative scaling of energy consumption with per capita output than other countries in this group, hence the high $\hat{\beta}_i$.

At the other end of the spectrum are countries with smaller coefficients. With a few exceptions, these are the countries with the largest GDPs per capita and growth driven primarily by low-intensity commercial sectors. Large energy consumers like the US are expected to have a high $\hat{\alpha}_i$, yet the sectoral shifts in such countries away from energy-intensive industry over the sample period drive negative signs on $\hat{\beta}_i$. Cuba is a notable exception. We speculate that Cuba's reduction in energy intensity stems from a gradual shift from an export-driven agricultural economy aimed at trade with the Soviet bloc to an economy dominated by tourism, a much less energy-intensive economic sector. Somewhat similarly, a change of emphasis away from heavy industry after Russia's economy was no longer dominated by the communist ideology of the Soviet system may explain Russia's reduction in energy intensity during this period.

Table 2 shows a subset of the results for $\hat{\alpha}_i$ and $\hat{\beta}_i$ estimated for the energy intensity-increasing group. In this group, where the surface shown in the right panel of Figure 2 generally increases but does not decrease, a positive $\hat{\beta}_i$ suggests an intensity increasing faster than that of the rest of the group, while a negative sign may suggest either a reduction in intensity or an intensity increasing more slowly that the rest of the group. For the most part, these countries were poorer over the sample period than those in the intensity-increasing group. Their economic growth tended be driven by agricultural or industrial exports. Hence, scaling up economic activity means scaling up the energy input in these sectors. Thus, the same logic that put South Korea and Hong Kong near the top of the previous list puts Thailand, Taiwan, and Vietnam near the top of this list: exports of manufactured goods and other energy intensive sectors.

The other countries are more difficult to explain. Afghanistan is clearly an outlier. The results for Afghanistan might be explained by decades of strife, so that both output and energy consumption are very low. Looking at the low end of the spectrum, six of them, Rwanda, Chad, Burkino Faso, Niger, Uzbekistan, and Burundi are landlocked countries, suggesting the possibility of relying on industries that do not require importing very much energy. But Afghanistan, Bolivia, and Malawi at the high end are also landlocked. A more detailed analysis of the industrial composition of each country might reveal further idiosyncrasies in endowments, geographical features, institutions, or historical events that are reflected in these coefficient estimates.

3.2 Elasticity Estimates

Although the functional coefficient estimates discussed above hold some interest in their own right, elasticity estimates are more economically meaningful and more widely applicable outside of the particular empirical application motivating our analysis. The left panel of Figure 3 shows the elasticity surface estimated for the intensity-reducing group. The exaggerated curvature to the
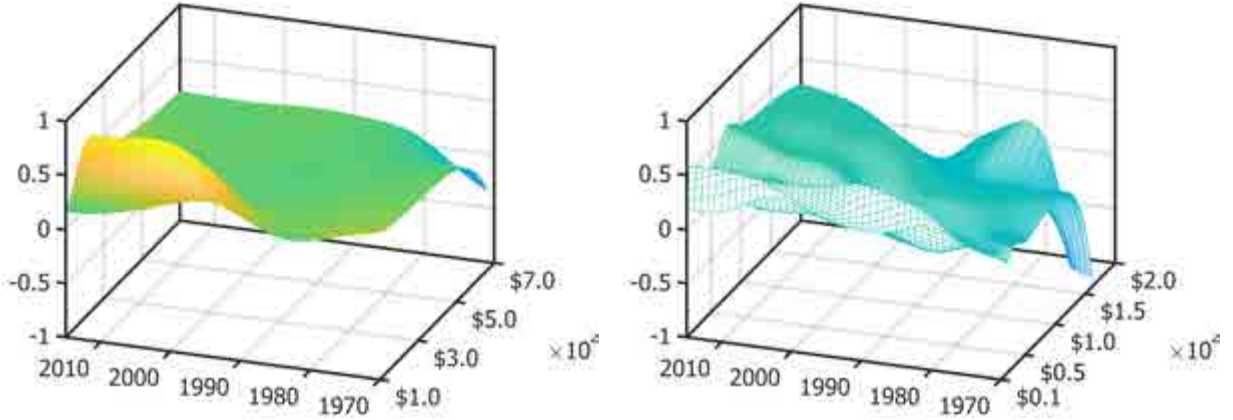
Figure 3: **Estimated Elasticity Surfaces.** Estimates of the common components $\Phi(r,x) = \Gamma(r,x) + \Gamma_x(r,x)x$ of the energy-GDP elasticities for the energy intensity-reducing group (left) and the energy intensity-increasing group (right). The surfaces are bias-corrected using the bootstrap procedure described in Appendix C.

extreme left and right of the surface give the impression of a nearly flat surface elsewhere, but although it is smooth this surface is in fact neither flat nor even monotonic. The right panel of the figure shows the elasticity surface estimated for the intensity-reducing group. We again discount the extreme lower left and upper right of the surface, but the middle of the surface is more obviously not flat in comparison with the surface estimated for the intensity-reducing group in the left panel.

The elasticity surfaces in Figure 3 are more resistant to ocular analysis than the coefficient surfaces in Figure 2. However, we see that the elasticity surfaces are approximately 0.2 to 0.3 for the intensity-reducing group and approximately -0.2 to 0.4 for the intensity-increasing group. Recall that these may be interpreted loosely as group average elasticities due to the identifying restriction in equation (5) that the idiosyncratic components of the slopes average to zero. Visual inspection gives the general impression that the latter group started with smaller or even negative average elasticities compared to the former group but ended up with larger average elasticities than the former group by the end of the sample, which is consistent with the negative correlation between growth rates of energy intensity and economic activity noted by Lescaroux (2011) and Csereklyei *et al.* (2016).

Early estimates such as those by Adams and Miovic (1968) and Brookes (1972) for developed countries tend to be high – even exceeding unity. Indeed, the period after the Second World War and up to when those two studies were published was characterized by very low energy prices that fueled energy-intensive economic growth for the wealthiest nations. It is well-documented that average elasticities have declined substantially since then, though the mechanisms for modeling the decline vary widely. A recent review by Liddle (2022) puts group averages across different studies
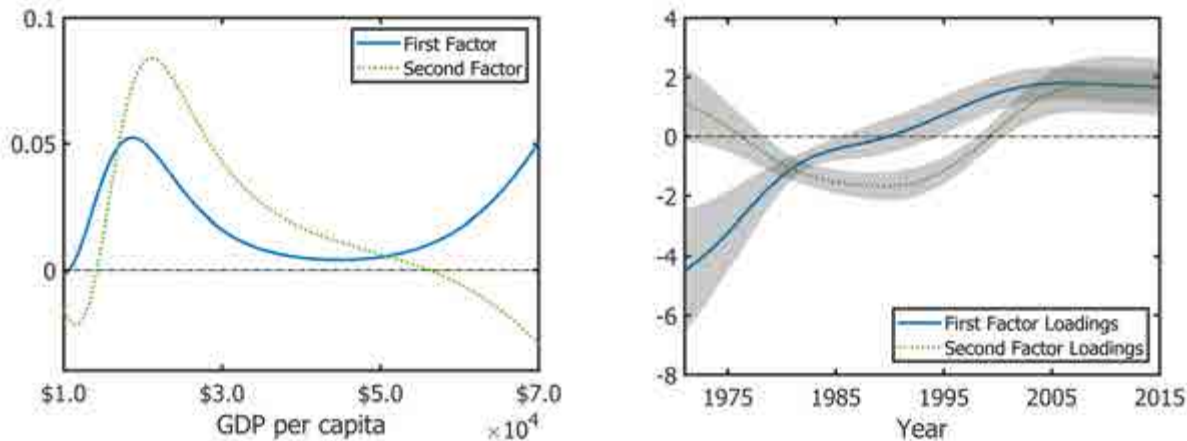
Figure 4: **Estimated FPCs and Loadings, Energy Intensity-Reducing Group.** Left: Estimated FPCs. Right: Estimated loadings. 95% confidence intervals created using the bootstrap procedure described in Appendix C.

in the range of 0.6 to 0.8 with the exception of an earlier study by the present authors, Chang *et al.* (2021). However, Csereklyei *et al.* (2016) note evidence for a decoupling of the relationship for developed countries, which is consistent with the results of Jakob *et al.* (2012) who find average elasticities for developed countries to be negative but statistically indistinguishable from zero. In the context of such a wide range of elasticities estimated in the extant literature, the group averages implied by our elasticity surface estimates are on the low end but certainly plausible.

To better understand the two surfaces estimated in Figure 3, we decompose each of them into FPCs and FPC loadings, as we explain in Section 2.3. The reader should keep in mind that both the factors and their loadings are defined for fluctuations of the functional time series of elasticities *around* its temporal average, so that a value of zero simply reflects that average. Figure 4 shows the leading two FPCs (left panel) against their respective loadings (right panel) for the elasticity surface of the energy intensity-reducing group. Together, these two FPCs and their loadings account for 88.6% of the total variation in the surface.

Both FPCs appear to have similar shapes for low-to-middle income countries in this group. Specifically, the FPCs are both relatively large over approximately $15,000-$55,000 per capita and have single peaks at approximately $20,000-$25,000. The first FPC attains a lower peak at a lower output level than the second FPC, but then the first peaks again at the upper end of the spectrum of economic activity. The main difference between the two FPCs occurs for high-income countries. Above approximately $55,000 per capita, the first FPC remains large relative to the temporal average elasticity and increases, while the second continues to decrease below the temporal average elasticity for these larger output levels.

Looking at the right panel of the figure, the loadings on these two FPCs begin with very different magnitudes at the beginning of the sample, but end up with similar magnitudes at the end. The loading on the first FPC increases nearly monotonically, and for this reason we dub it the *intensity-increasing factor*. The time-varying rate at which the intensity-increasing factor drives intensity shown by the loading is roughly consistent with a history of energy prices. The loading increases rapidly during the 1970s, then it slows down during the 1980s, then it increases rapidly again during the 1990s, and then finally it levels out and perhaps even decreases slightly after the turn of the century.

In contrast, the loading on the second FPC appears to have a "U" shape until about the turn of the century, falling below the temporal average in the late 1970s and rising above it in the late 1990s. After about 2005, the loading on the second FPC levels out into a pattern almost identical to that of the first FPC. Because of the loading clearly declines over much of the sample, we dub this FPC the *intensity-reducing factor*. While the name may seem paradoxical in light of clear increase in its loading over the period of about 1990-2005, this was also a period of global prosperity, so that the decrease in the factor itself with output could have negated much of that increase for countries in this group.

Looking again at the loading on the intensity-increasing factor, the leveling out and perhaps slight decrease of the loading on the energy intensity factor after the turn of the century is consistent with a reduction in energy intensity that earns this group of countries its intensity-reducing appellation in addition to the 1970s and 1980s. As hinted at above, the slowdown could be indicative of changes in technology and preferences in an era of generally high and unpredictable energy prices, which spur more measured use of energy. Alternatively or additionally, the slowdown could be indicative of an increasing awareness of the effects of fossil fuel consumption on local environments and the global climate and consequent effects on human health and the economy, though one could argue that it has been much more recently – if ever – that countries have made a serious effort to curtail emissions other than sulfur dioxide and CFCs. The slowdown in the loading on the intensity-increasing factor could be indicative of sectoral shifts away from manufacturing and other energy-intensive industries in the wealthier developed countries, for which the intensity-increasing factor is relatively large. However, the timing of this pattern of sectoral shifts seems to be explained better by the intensity-reducing factor.

Looking only at the period after the mid-1980s, the loading on the intensity-reducing factor has more of an "S" shape than the "U" shape noted above. Specifically, the most rapid increase in the loading occurs over approximately 1990-2005. Given that the second FPC is decreasing in economic activity, this shape suggest the possibility that wealthier countries, for which this factor has a low weight, exported energy-intensive industries over this period to middle-income countries also in this group, like South Korea and Mexico, for which this factor has a high weight, or to
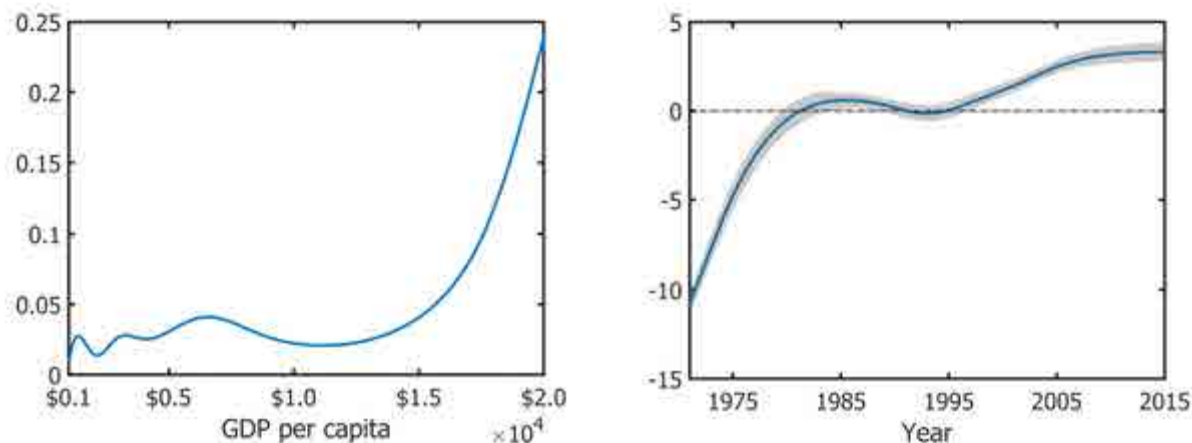
Figure 5: **Estimated FPCs and Loadings, Energy Intensity-Increasing Group.** Left: Estimated FPCs. Right: Estimated loadings. 95% confidence intervals created using the bootstrap procedure described in Appendix C.

middle-income or low-income countries in the intensity-increasing group, like Taiwan and China. Indeed the period since the mid-1980s witnessed unprecedented global trade integration sufficient to make such shifts possible.

Technology, preferences, and prices could also play a role in shaping the intensity-reducing factor. The unexpectedly high prices of the 1970s and 1980s could have contributed to a decrease in the load on the intensity-reducing factor just as they may have contributed to the slowing of the increase of the load on the intensity-increasing factor. Environmental movements began to gain traction with legislation during this period, such as the Clean Air Act (1970) and its amendments (1977 and 1990) in the US.

Moving now to the energy intensity-increasing group and revisiting Figure 3, the right panel shows the elasticity surfaces estimated for the second group. As we mention above, it is difficult to make generalizations from the figure, but it appears that for medium-income countries – i.e., the wealthiest ones in this group – a peak occurred in the mid-1980s, followed by a downturn and then by a more recent increase. Such an increase is qualitatively similar to the pattern that Adams and Miovic (1968) and Brookes (1972) observe for developed countries prior to our sample. In that light, this result echoes that of van Benthem (2015), who finds the benefit of technologies that can reduce energy intensity in developing economies is offset by more energy-intensive consumption bundles.

Figure 5 shows the leading FPC (left panel) and its loading (right panel) for the intensity-increasing group, which alone explain 83.9% of the variation in the surface. The FPC is positive for all levels of output. It remains relatively flat up to about $15,000$, after which it increases

precipitously to the maximum output at which it is estimated, $20,000 per capita. In fact, from $10,000 to $20,000 this FPC strongly resembles the leading FPC in the intensity-reducing group up to an elongation of the vertical axis.

So, too, does this FPC's loading resemble that of the leading FPC of the intensity-reducing group. It increases dramatically during the 1970's, slowing down with the oil prices hikes and recessions of the late 1970s and early 1980s. Like the the loading on the first FPC of the intensity-decreasing group, it accelerates again in the 1990s and levels out in the early 2000s. It is no surprise, then, that we label this the *intensity-increasing factor* like the first FPC of the intensity-reducing group.

The combined effect of the energy intensity factor and its loading suggest a strong positive effect on elasticities as output increases over time. Indeed, this is consistent with export-driven industrial growth that many of the wealthiest countries in this group have used to fuel their development. While both groups have an intensity-increasing factor that generally increases over time, this group does not have the intensity-reducing factor that mitigates the increase for the wealthiest countries in the intensity-reducing group.

The intensity-increasing factors given by the leading FPCs are similar across both groups. These generally reflect the global buildup of energy-intensive sectors over the twentieth century, but this buildup has stopped and perhaps declined since the turn of the century. This is not to say that economic growth has slowed nor even that economic growth in energy-intensive sectors has slowed, but simply that the proportion of energy-intensive sectors to all sectors has stopped increasing. In this context and to the extent that the intensity-reducing factor reflects sectoral shifts, it layers on a redistribution of energy-intensive sectors away from the wealthiest countries and onto middle-income and low-income countries.

3.3   Country-Specific Elasticities

We have alluded both to country-specific idiosyncrasies in discussing the functional coefficient surface estimates and to the difficulty in interpreting FPCs and their loadings beyond a snapshot of the common component of elasticities for a fixed level of real GDP or fixed time period. But a large part of the appeal of our panel approach is the ability to identify elasticities when both arguments are correlated and changing over time. As with the coefficient surfaces, countries tend to move from lower right to upper left of the elasticity surfaces as their economies grow over time. For this reason, it is possible that elasticities of specific countries in the intensity-increasing group do not increase very much, while those in the intensity-reducing group may not decrease very much if at all, even setting aside the idiosyncratic components of each country's elasticity.

In this light, Figure 6 traces out the estimated energy-GDP elasticities of eight large economies along the elasticity surface corresponding to each country's group and using estimates of the idiosyn-
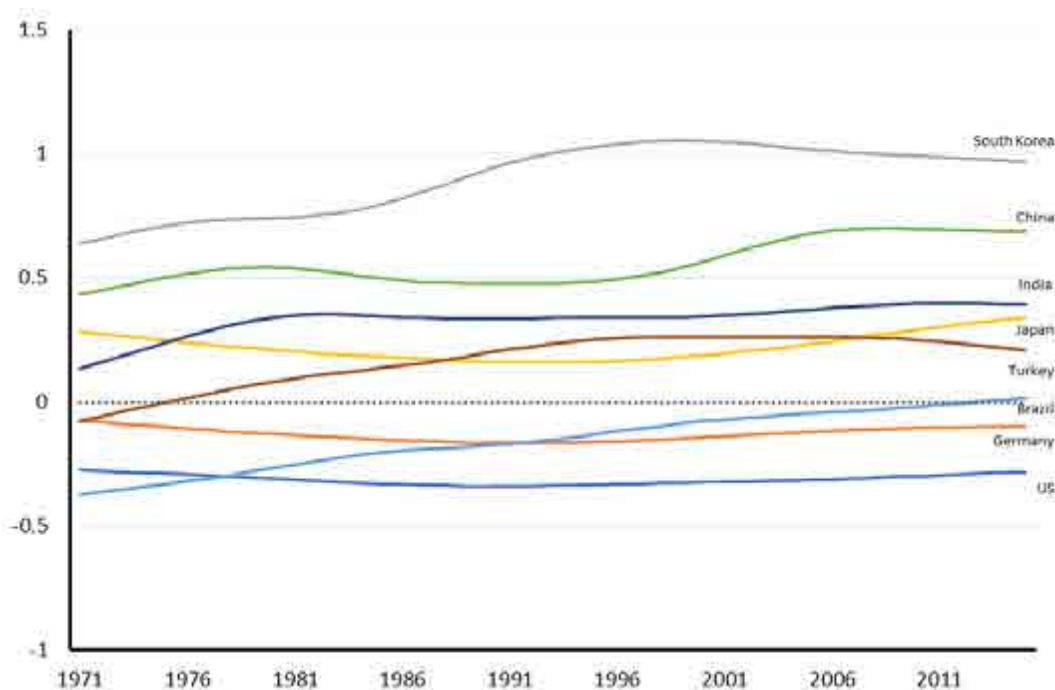
Figure 6: **Estimated Elasticities of Selected Countries.** Estimated elasticities of four countries in the energy intensity-reducing group (US, Germany, South Korea, Japan) and four countries in the energy intensity-increasing group (Brazil, China, India, Turkey).

cratic component β_i of the slope estimated for each country. Specifically, elasticity estimates are plotted for the US, Germany, South Korea, and Japan (countries in the energy intensity-reducing group), and for Brazil, China, India, and Turkey (countries in the energy intensity-increasing group).

Of the eight countries examined, South Korea is the most energy intensive for the duration of the sample, with near unit elasticity since the early 1990s. A member of the intensity-reducing group, South Korea appears to have a declining elasticity since the turn of the century, but this decline is very shallow compared to the strong increase in its elasticity during the 1970s and 1980s. South Korea's trajectory is easy to interpret, because it started as a relatively poor country and boosted its energy-intensive industrial base so that its energy intensity by the early 1990s resembled that of Western Europe decades previously (Adams and Miovic, 1968).

The second-most energy intensive is China. China also began as a relatively poor country and followed an industrial export-based model like that of South Korea. Indeed, we see a similar pattern in that China's elasticity starts low and relatively stable, but then increases from the mid-1990s to the mid-2000s, after which it levels out. Roughly speaking, the pattern of the Chinese elasticity appears to lag that of South Korea by a decade or two, though up to a vertical shift because it

never exceeds 0.8.

The pattern of Turkey's elasticity appears somewhat similar to that of South Korea in that it increases until about the turn of the century after which it shows a noticeable decline. Indeed, one could attribute the inverted "U" to an EKC, but we believe this is an unwarranted oversimplification. The time period – but not necessarily output level – over which Turkey's elasticity declines is roughly consistent with the declining elasticities of other OECD countries, even though Turkey is one of the few OECD countries in the intensity-increasing group (See Appendix A.2 for more details).

Like South Korea, China, and Turkey, India's elasticity increased during the 1970s but has been stable at about 0.4 since then. In contrast, Brazil's elasticity has never stopped increasing. However, Brazil's elasticity was negative at the beginning of the sample and only recently reached zero.

Germany, Japan, and the US have similar patterns in that their elasticities declined rather than increased until the 1990s. The elasticities of Germany and the US have remained nearly constant and negative since then, while Japan's has slightly increased. Given that these three are among the wealthiest countries in the sample in terms of real GDP per capita and setting aside the obvious correlation of time and output, these patterns are inconsistent with an EKC unless the peak of the EKC happened prior to the 1970s, which seems unlikely. A simpler and more plausible explanation for the decline in the elasticities during the 1970s is a response of these advanced economies to the energy price hikes of that decade.

The negative sign of elasticities observed for Brazil, Germany, and the US seems surprising at first glance, because the energy-GDP elasticity that we examine is easily confused with the income elasticity of energy. A negative income elasticity would mean that energy is an inferior good. It may very well be the case that consumers view certain types of energy, such as those that generate high amounts of sulfur dioxide that may result in acid rain, as inferior goods. But it is hard to imagine that energy from all sources would be an inferior good at a macroeconomic scale, especially for a country like the US that is a large polluter.

It is important to keep in mind our discussion at the beginning of Section 3 of the energy-GDP elasticity – it is not an income elasticity and would be difficult to interpret as such without including additional covariates to control for energy prices and both supply and demand shifters. Instead, a negative value simply indicates an AEEI, which is an increase in efficiency. In other words, Brazil, Germany, and the US have grown in spite of decreasing energy consumption, essentially decoupling economic growth from energy consumption, while the other five countries in the Figure 6 have increased energy consumption in order to grow, albeit the increases have happened at a slower rate than economic growth because the elasticities do not exceed unity. This is less a statement about technical efficiency than it is about the allocation of economic activity across sectors with different

energy intensities, as the example of Iceland discussed previously illustrates.

Our results on decoupling largely agree with those of Csereklyei *et al.* (2016). While decoupling is observed in some countries, and especially in developed countries, it is not prevalent enough to affect the common components of the elasticities of each group, which are clearly positive over most values of (r, x) as we observed from Figure 3. However, given some of the large negative values of β_i in Tables 1 and 2, the phenomenon appears to be more prevalent than previously understood using less flexible models.

4   Conclusion

In this paper, we embrace nonlinearity and both individual and temporal heterogeneity in the long-run relationship between aggregate economic activity and energy consumption. Our methodological contribution is based on semiparametric estimation of a functional coefficient panel model, which allows heterogeneity in an economic relationship at the same time as a common component that is both nonlinear and varies over time.

Using this approach, we can effectively distinguish nonlinearity in the covariate from a time-varying linear or otherwise parametrically specified relationship. We emphasize the unique ability of a nonparametric functional panel approach in identifying a functional relationship with two arguments. Further, we use functional principal components analysis to examine commonalities across across the domain of the covariate (the FPCs or factors) and across time (the loadings), and we introduce a novel bootstrap procedure for inference and interval estimates on the loadings. Our use of functional principal component analysis to extract common FPCs in a relationship of interest using panel data is new to the best of our knowledge.

We specifically examine national energy-GDP elasticities, the derivative of log energy consumption with respect to log real GDP for a specific country, which is a key elasticity in understanding the efficiency and evolution of aggregate usage of energy to produce output and a critical input that economists use to model the interaction of the climate and economic systems using integrated assessment models. However, our methodology is quite general and can be applied to any panel with reasonably large temporal and cross-sectional dimensions.

Our main results suggest that developed countries – those in the energy intensity-reducing group – generally have had declining elasticities in recent years with a common component that has fluctuated within a range of about 0.2 to 0.3 over the sample. On the other hand, relatively underdeveloped countries – those in the energy intensity-increasing group – share a pattern of elasticities that is generally increasing over time and similar to that followed by developed countries in previous decades but fluctuation within a range of about -0.2 to 0.4 over the sample. Our innovation lies not only in generating these estimates by way of a semiparametric functional coefficient

model, but also in sorting these countries into their respective groups based on common patterns in the functional coefficients.

Indeed, all countries appear to share an energy-increasing factor that echoes major changes in global energy prices, while the more developed countries also have an intensity-declining factor that may result from sectoral shifts, energy prices, technology, and/or energy sources preferences and that offsets the growing intensity and has even decoupled economic growth from energy consumption for the wealthiest countries, such as Germany and the US. Decoupling of this relationship has been noted in the extant literature, but our results suggest that it may be more prevalent than is currently thought.

References

Adams F.G. and P. Miovic (1968). On relative fuel efficiency and the output elasticity of energy consumption in Western Europe, The Journal of Industrial Economics 17, 41-56.

Ahn, S.C. and A.R. Horenstein (2013). Eigenvalue ratio test for the number of factors, Econometrica 81, 1203-1227.

Arisoy, I. and I. Ozturk (2014). Estimating industrial and residential electricity demand in Turkey: A time varying parameter approach, Energy 66, 959-964.

Bai, J. and S. Ng (2004). A PANIC attack on unit roots and cointegration, Econometrica 72, 1127-1177.

van Benthem, A.A. (2015). Energy leapfrogging, Journal of the Association of Environmental and Resource Economists 2, 93-132.

Bierens, H.J. and L.F. Martins (2010). Time-varying cointegration, Econometric Theory 26, 1453-1490.

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.

Brookes, L.G. (1972). More on the output elasticity of energy consumption, Journal of Industrial Economics 21, 83-94.

Burke, P.J. and Z. Csereklyei (2016). Understanding the energy-GDP elasticity: A sectoral approach, Energy Economics 58, 199-210.

Cai, Z. (2007). Trending time-varying coefficient time series models with serially correlated errors, Journal of Econometrics 136, 163-188.

Cai, Z., Q. Li and J.Y. Park (2009). Functional-coefficient models for nonstationary time series data, Journal of Econometrics 148, 101-113.

Chamberlain, G. and M. Rothschild (1983). Arbitrage, factor structure and mean-variance analysis in large asset markets, Econometrica 51, 1305-1324.

Chang, Y., M. Choi and J.Y. Park (2023). A factor model for functional time series, mimeograph.

Chang, Y., Y. Choi, C.S. Kim, J.I. Miller, 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, J.Y. Park (2021). Forecasting regional long-run energy demand: A functional coefficient panel approach, Energy Economics 96, 105117.

Chang, Y., R.K. Kaufmann, C.S. Kim, J.I. Miller, J.Y. Park, and S. Park (2020). Evaluating trends in time series of distributions: A spatial fingerprint of human effects on climate, Journal of Econometrics 214, 274-294.

Chang, Y., C.S. Kim, and J.Y. Park (2016b). Nonstationarity in time series of state densities, Journal of Econometrics 192, 152-167.

Chen, Y.-H.H., S. Paltsev, A. Gurgel, J.M. Reilly, and J. Morris (2022). The MIT EPPA7: A multisectoral dynamic model for energy, economic, and climate scenario analysis. Joint Program Report Series Report 360.

Cooley, T.F. and E.C. Prescott (1976). Estimation in the presence of stochastic parameter variation, Econometrica 44, 167-184.

Csereklyei, Z., M.d.M. Rubio Varas, and D.I. Stern (2016). Energy and economic growth: The stylized facts, 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.

Inglesi-Lotz, R. (2011). The evolution of price elasticity of electricity demand in South Africa: A Kalman filter application, Energy Policy 39, 3690-3696.

Jakob, M., M. Haller, and R. Marschinski (2012). Will history repeat itself? Economic convergence and convergence in energy use patterns, Energy Economics 34, 95-104.

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.

Kraft, J. and A. Kraft (1978). On the relationship between energy and GDP, The Journal of Energy and Development 3, 401-403.

Lescaroux, F. (2011). Dynamics of final sectoral energy demand and aggregate energy intensity, Energy Policy 39, 66-82.

Liddle, B. (2022). What is the temporal path of the GDP elasticity of energy consumption in OECD countries? An assessment of previous findings and new evidence, Energies 15, 3802.

Liddle, B. and H. Huntington (2020). Revisiting the income elasticity of energy consumption: A heterogeneous, common factor, dynamic OECD & non-OECD country panel analysis, The Energy Journal 41, 207-229.

Liddle, B., R. Smyth, and X. Zhang (2020). Time-varying income and price elasticities for energy demand: Evidence from a middle-income panel, Energy Economics 86, 104681.

Lin, C.-F.J., and T. Teräsvirta (1994). Testing the constancy of regression parameters against continuous structural change, Journal of Econometrics 62, 211-228.

Medlock, K.B. and R. Soligo (2001). Economic development and end-use energy demand, The Energy Journal 22, 77-105.

Newbold, P. and T. Bos (1985). Stochastic Parameter Regression Models. Sage Publications, Newbury Park, CA.

Paltsev, S., J.M. Reilly, H.D. Jacoby, R.S. Eckaus, J. McFarland, M. Sarofim, M. Asadoorian, and M. Babiker (2005). The MIT Emissions Prediction and Policy Analysis (EPPA) Model: Version 4. Joint Program Report Series Report 125.

Park, J.Y. and S.B. Hahn (1999). Cointegrating regressions with time varying coefficients, Econometric Theory 15, 664-703.

Park, S.Y. and G. Zhao (2010). An estimation of U.S. gasoline demand: A smooth time-varying cointegration approach, Energy Economics 32, 110-120.

Ramsay, J.O. and B.W. Silverman (2005). Functional Data Analysis, in: Springer Series in Statistics, Second Edition. Springer, New York.

Richmond, A.K. and R.K. Kaufmann (2006). Is there a turning point in the relationship between income and energy use and/or carbon emissions? Ecological Economics 56, 176-189.

Sargent, T.J., and C.A. Sims (1977). Business cycle modeling without pretending to have too much a priori economic theory, in C.A. Sims (ed.), New Methods in Business Research. Federal Reserve Bank of Minneapolis, Minneapolis.

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.

A   Appendix: Details of the Data

A.1   Missing Data

Estimating the model requires only two variables, energy consumption per capita and GDP per capita. However acquiring these variables for a large panel is not straightforward. In fact, we need not only energy consumption and GDP, but we also need population in order to create per capita observations. The final data set is an unbalanced panel consisting of 180 countries over 1971-2015. The population of these 180 countries covers 91.1% of the UN estimated world population in 1971 and 98.6% by 2015. 30 of these countries do not have data over the entire span, and the largest increase in coverage, 5.3%, occurs in 1996 with the addition of former Soviet bloc countries. Details of how we construct this panel follow.

Energy Consumption.   Energy consumption is measured in millions of tonnes of oil equivalent (Mtoe), and is available from Enerdata over durations up to 1970-2016 for 186 countries. However, we drop data for 1970 and 2016, because data for only 71 and 48 countries are available in those respective years at the time we conducted our analysis. We omit data for most former Soviet bloc countries in Eastern Europe until 1995. All data for Eritrea, North Korea, Somalia, Kosovo, Montenegro, and Netherlands Antilles are dropped, either because not enough energy consumption data are available or because not enough GDP data are available. Energy consumption for Niger for 1979 is missing.

Population.   Population data are available from the World Bank for 217 countries over durations up to 1960-2015. Population data are also available from Enerdata and the Penn World Table (PWT), but we defer to the World Bank data for countries and years for which such data are available. Substantial discrepancies between the three sources are only evident for Sudan. We augment the data with Taiwanese population data from Enerdata, because Taiwan is not a member of the World Bank. Population data for Kuwait are missing for 1992-94. Data on population prior to 1971 for all countries are dropped.

GDP.   GDP data are available mainly from the PWT version 9.0 for 182 countries over durations up to 1950-2014. We employ Enerdata and World Bank GDP data where available for countries for which energy consumption data are available. PWT real GDP data are chained to 2011 (output-side GDP at chained PPPs in millions of 2011 US$), while those from Enerdata and the World Bank are chained to 2005. Where the latter data sources are used, we match base years. The PWT GDP data to 2014 from are augmented by data from the World Bank and Enerdata for 2015. We also use World Bank and Enerdata to fill data that are available from those sources but missing

from the PWT. For most (173) countries, we use World Bank GDP growth rates from 2014-2015 to estimate PWT GDP in 2015.

For Taiwan and Syria (2), World Bank GDP data are not available, but we use Enerdata growth rates instead. For Cuba, Guyana, Kiribati, Papua New Guinea, Tonga, Vanuatu, and Samoa (7), no PWT GDP data are available, but World Bank data are available. We use nominal GDP in 2011 to capture the difference in base years and the growth rates in real GDP to reconstruct an estimate of real GDP chained to 2011 like that of the PWT. For Afghanistan, Libya, and Solomon Islands (3), no PWT GDP data are available, some World Bank data are available, but Enerdata data are available. The procedure for these three countries is the same for the latter seven, but Enerdata growth rates are used where World Bank data are unavailable.

A.2 Country Groups

The 180 countries in our sample exhibit coefficient heterogeneity not only in β_i but also in $\Gamma(r, x)$. Roughly speaking, developing countries appear to have $\Gamma(r, x)$ increasing in t, while developed countries have $\Gamma(r, x)$ increasing in r only up to about 2004, after which point the function declines for many values of x. In order to deal with this heterogeneity systematically, we first break the countries into two groups: OECD (35 countries) and non-OECD (145 countries) and then subsequently refine the groups by adding or subtracting countries to or from each group. We now describe the procedure used to accomplish this grouping.

As the starting point, we fit the model in equation (4) separately for OECD and non-OECD groups to obtain the surfaces $\hat{\Gamma}_{OECD}(r, x)$ and $\hat{\Gamma}_{NonOECD}(r, x)$ for each group. Next, for each country i, regardless of its group, we estimate

$$Y_{it} - \hat{\Gamma}^*_{OECD}(t, X_{it})X_{it} = \alpha_i + \beta_{*i}X_{it} + e_{it}$$
$$Y_{it} - \hat{\Gamma}^*_{NonOECD}(t, X_{it})X_{it} = \alpha_i + \beta_{*i}X_{it} + e_{it}$$

using these surfaces.

We then compare the respective root mean squared errors (RMSEs) $RMSE_i^{OECD}$ and $RMSE_i^{NonOECD}$ from these two regressions. Any OECD country for which $RMSE_i^{OECD} > RMSE_i^{NonOECD}$ is removed from the OECD group. There are two, Latvia and Turkey, while 33 remain. Any non-OECD country for which $x_{i2015} > \log(15, 698)$ and $RMSE_i^{OECD} < RMSE_i^{NonOECD}$ is removed from the non-OECD group. There are 14, Bahrain, Bahamas, Belarus, Bermuda, Bulgaria, Cuba, Cyprus, Hong Kong, Kuwait, Lithuania, Panama, Romania, Russia, and Uruguay, while 131 remain.

In which group do these 16 removed countries belong? There are $2^{16} = 65, 536$ configurations in which these 16 countries belong to one group or another. We re-estimate the model for each group using each of these 2^{16} groupings, retaining the fitted residuals for each country and year and
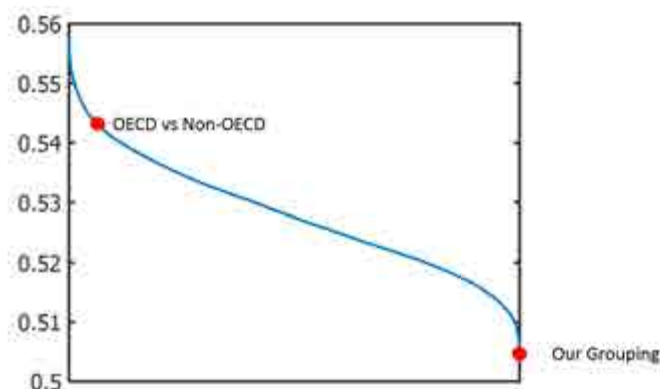
Figure A.1: **Root Mean Squared Errors.** RMSEs calculated for $2^{16} = 65,326$ groupings with RMSE of original OECD vs. non-OECD groups and optimized RMSE of intensity-decreasing vs. intensity-increasing groups indicated.

calculating a population-weighted root mean squared error for the whole panel. We then pick the combination that minimizes root mean squared error across all 2^{16} groupings. Figure A.1 shows the clear improvement in model fit resulting from assigning countries to groups in this way relative to the OECD/non-OECD starting point.

Based on this procedure, Turkey moves from the OECD group to the non-OECD group, while Latvia and the 33 original countries remain. As explained in the text, we label this the energy intensity-reducing group. On the other hand, 12 non-OECD countries, Bahrain, Belarus, Bermuda, Bulgaria, Cuba, Cyprus, Hong Kong, Kuwait, Lithuania, Panama, Romania, Russia move to the OECD group, while Bahamas, Uruguay, and the original 131 remain, and we label this the energy intensity-increasing group. The specific countries in each group are listed below.

Energy Intensity-Reducing Group (46 Countries): 35 OECD countries, minus 1, plus 12 marked by a + superscript, ordered by ISO three-letter country codes. Australia, Austria, Bahrain[+], Belarus[+], Belgium, Bermuda[+], Bulgaria[+], Canada, Chile, Cuba[+], Cyprus[+], Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hong Kong[+], Hungary, Iceland, Ireland, Israel, Italy, Japan, Kuwait[+], Latvia, Lithuania[+], Luxembourg, Mexico, Netherlands, New Zealand, Norway, Panama[+], Poland, Portugal, Romania[+], Russia[+], Slovakia, Slovenia, South Korea, Spain, Sweden, Switzerland, United Kingdom, United States.

Energy Intensity-Increasing Group (134 Countries): 145 non-OECD countries, minus 12, plus 1 marked by a + superscript, ordered by ISO three-letter country codes. Afghanistan, Albania, Algeria, Angola, Argentina, Armenia, Azerbaijan, Bahamas, Bangladesh, Barbados, Belize, Benin,

Bhutan, Bolivia, Bosnia and Herzegovina, Botswana, Brazil, Brunei, Burkina Faso, Burundi, Cabo Verde, Cambodia, Cameroon, Central African Republic, Chad, China, Colombia, Comoros, Congo (DR), Congo (R), Costa Rica, Côte d'Ivoire, Croatia, Djibouti, Dominica, Dominican Republic, Ecuador, Egypt, El Salvador, Equatorial Guinea, Ethiopia, Fiji, Gabon, Gambia, Georgia, Ghana, Grenada, Guatemala, Guinea, Guinea-Bissau, Guyana, Haiti, Honduras, India, Indonesia, Iran, Iraq, Jamaica, Jordan, Kazakhstan, Kenya, Kiribati, Kyrgyzstan, Laos, Lebanon, Lesotho, Liberia, Libya, Macao, Macedonia, Madagascar, Malawi, Malaysia, Maldives, Mali, Malta, Mauritania, Mauritius, Moldova, Mongolia, Morocco, Mozambique, Myanmar, Namibia, Nepal, Nicaragua, Niger, Nigeria, Oman, Pakistan, Papua New Guinea, Paraguay, Peru, Philippines, Qatar, Rwanda, Saint Lucia, Saint Vincent and the Grenadines, Samoa, Sao Tome and Principe, Saudi Arabia, Senegal, Serbia, Seychelles, Sierra Leone, Singapore, Solomon Islands, South Africa, Sri Lanka, Sudan, Suriname, Swaziland, Syria, Taiwan, Tajikistan, Tanzania, Thailand, Togo, Tonga, Trinidad and Tobago, Tunisia, Turkey[+], Turkmenistan, Uganda, Ukraine, United Arab Emirates, Uruguay, Uzbekistan, Vanuatu, Venezuela, Vietnam, Yemen, Zambia, Zimbabwe.

B   Appendix: Additional Details of Estimation

B.1   Bandwidth Selection

There is no theory or empirical consensus guiding the choice of bandwidths h_r and h_x with a nonstationary panel. We use the formulas $h_r = cn^{-1/6}$ and $h_r = c\tau n^{-1/6}$ where τ is the ratio of the standard deviation of t/T, all time observations normalized to the unit interval, to that of $X_{it}/\max(X_{it})$, all log real GDP observations normalized to the unit interval. Specifically, $\tau = 0.261$ for the energy intensity-reducing group and $\tau = 0.367$ for the energy intensity-reducing group. We use wider bandwidths of $h_r = cn^{-1/8}$ and $h_r = c\tau n^{-1/8}$ for estimating the derivatives.

We select the bandwidth constant c based on out-of-sample forecast performance. Specifically, we compare RMSEs of out-of-sample forecasts with the last 5 years or 10 years of the sample reserved for this purpose. Because the forecasts for choosing c have relatively short horizons, the function $\Gamma(t, X_{it})$ is forecast simply by extending the standard normal kernel up to the end of the sample, $t = 2015$ and $X_{i,2015}$. More precise forecasting may be possible if forecasting were the goal, but our goal here is simply bandwidth selection.

For the OECD group and the similar intensity-reducing group, we choose a bandwidth constant of $c = 0.40$ based on bandwidth constants of 0.20 and 0.34 selected for 5-year ahead and 10-year ahead forecasts. The RMSE at 0.4 is barely larger than that at 0.34. For the non-OECD group and similar intensity-increasing group, we choose a bandwidth constant of $c = 0.30$ based on bandwidth constants of 0.24 and 0.20 selected for 5-year ahead and 10-year ahead forecasts. Deliberately setting c to be a bit high imposes smoothness on the estimates, which is particularly useful for the

second group, which otherwise shows a considerably less smooth surface estimate $\hat{\Gamma}(r, x)$ than that of the first group.

B.2 Correlation of α_i with β_i

As noted in the paper, we observe a strong negative correlation between our estimates of α_i and β_i. To understand why, rewrite our model in (4) as

$$Z_{it} = \alpha_i + \beta_i X_{it} + U_{it}$$

by defining $Z_{it} = Y_{it} - \Gamma(t, X_{it})X_{it}$. The term α_i may be decomposed as

$$
\begin{aligned}
\alpha_i &= \bar{Z}_i - \beta_i \bar{X}_i - \bar{U}_i \\
&= [\bar{Z} - \beta_i \bar{X}] + [(\bar{Z}_i - \bar{Z}) - \beta_i(\bar{X}_i - \bar{X}) - \bar{U}_i] \\
&= \alpha_{1i} + \alpha_{2i},
\end{aligned}
$$

where $\bar{P}_i = (1/T)\sum_{t=1}^{T} P_{it}$ for $P = Z, X, U$, and $\bar{Q} = (1/NT)\sum_{i=1}^{N}\sum_{t=1}^{T} Q_{it}$ for $Q = Z, X$.

Define

$$\sigma(\alpha, \beta) = \frac{1}{N}\sum_{i=1}^{N}(\alpha_i - \bar{\alpha})(\beta_i - \bar{\beta}),$$

$\sigma^2(\alpha) = \sigma(\alpha, \alpha)$, and $\sigma^2(\beta) = \sigma(\beta, \beta)$ to be the sample covariance and variances of (α_i) and (β_i). The sample correlation is therefore given by $\rho(\alpha, \beta) = \sigma(\alpha, \beta)/[\sigma(\alpha)\sigma(\beta)]$. Define $\sigma(\alpha_k, \beta)$, $\sigma^2(\alpha_k)$ and $\rho(\alpha_k, \beta)$ similarly for $k = 1$ and 2.

For our model, $(\bar{Z}_i - \bar{Z})$ and $(\bar{X}_i - \bar{X})$ are negligible compared to $\bar{X}$, and therefore, $\alpha_i \approx \alpha_{1i}$ so that $\sigma(\alpha_1)/\sigma(\alpha) \approx 1$. Consequently,

$$\rho(\alpha, \beta) = \rho(\alpha_1, \beta)\left(1 + \frac{\sigma(\alpha_2, \beta)}{\sigma(\alpha_1, \beta)}\right)\frac{\sigma(\alpha_1)}{\sigma(\alpha)} = -\left(1 + \frac{\sigma(\alpha_2, \beta)}{\sigma(\alpha_1, \beta)}\right)\frac{\sigma(\alpha_1)}{\sigma(\alpha)} \approx -1.$$

The ratios above can be computed from the estimates $(\hat{\beta}_i)$ of (β_i) and $(\hat{\bar{U}}_i)$ of $(\bar{U}_i)$. Using our data, these ratios are given by

$$\frac{\sigma(\alpha_1)}{\sigma(\alpha)} = 0.905 \quad \text{and} \quad \frac{\sigma(\alpha_2, \beta)}{\sigma(\alpha_1, \beta)} = 0.098,$$

which implies that $\rho(\alpha, \beta) = -0.993$. In other words, the estimates $(\hat{\alpha}_i)$ have an almost perfect negative correlation with $(\hat{\beta}_i)$, a result that reflects correlation in the underlying heterogeneity, as we have just shown, rather than correlations from estimation error or a spurious regression.

C   Appendix: Bootstrap Procedure

We utilize a bootstrap procedure to compute standard errors and confidence intervals to assess uncertainty. Moreover, bandwidths of nonparametric estimators inherently balance bias and variance, so the bootstrap may further reduce the bias.

C.1   Construction

The discussion in this section follows that of a closely related paper by the authors (Chang *et al.*, 2021, Appendix A.1). However, there are some differences in notations and groups. For either the intensity-reducing or the intensity-increasing group, let N denote the number of countries for which we observe data over the whole time span, 1971-2015, so that $T = 45$. Bootstrapped residuals used to calculate interval estimates for the parameters and FPC loadings are obtained using the following steps.

1. Obtain fitted residuals $\hat{U}_{it}$ from the regression in (8), and temporally demean them to obtain $\hat{u}_{it} = \hat{U}_{it} - \overline{\hat{U}_i}$.

2. Calculate the principal components of the covariance matrix of $\hat{u}_{\cdot t} = (\hat{u}_{1t}, ..., \hat{u}_{Nt})'$. Decompose $\hat{u}_{\cdot t}$ as $\hat{u}_{\cdot t} = \hat{\Lambda}\hat{g}_t + \hat{\eta}_{\cdot t}$, where $\hat{\Lambda} = (\hat{\lambda}_1, ..., \hat{\lambda}_N)'$ with $\hat{\lambda}_i$ defined to be an $r \times 1$ vector of factors for variable $(\hat{u}_{i\cdot})$ with an $r \times 1$ vector $\hat{g}_t$ of factor loadings, and where $\hat{\eta}_{\cdot t} = (\hat{\eta}_{1t}, ..., \hat{\eta}_{Nt})'$ is a vector of idiosyncratic components of $\hat{u}_{\cdot t}$. The number of factors r is chosen by agreement between the eigenvalue ratio and growth ratio tests (Ahn and Horenstein, 2013).

3. Model $\hat{g}_t$ as a vector autoregression,

$$\hat{g}_t = B_1\hat{g}_{t-1} + \cdots B_k\hat{g}_{t-k} + \epsilon_t \tag{C.1}$$

with the autoregressive order k determined by BIC. Denote parameter estimates and fitted residuals by $\hat{B}_1, ..., \hat{B}_k$ and $\hat{\epsilon}_t$ respectively. Re-sample from $(\hat{\epsilon}_t - \overline{\hat{\epsilon}}_\cdot)$ to obtain bootstrap samples (ϵ_t^*), from which bootstrap samples (g_t^*) are obtained using

$$g_t^* = \hat{B}_1 g_{t-1}^* + \cdots \hat{B}_k g_{t-k}^* + \epsilon_t^*, \tag{C.2}$$

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

4. Model $\hat{\eta}_{it}$ as an autoregression,

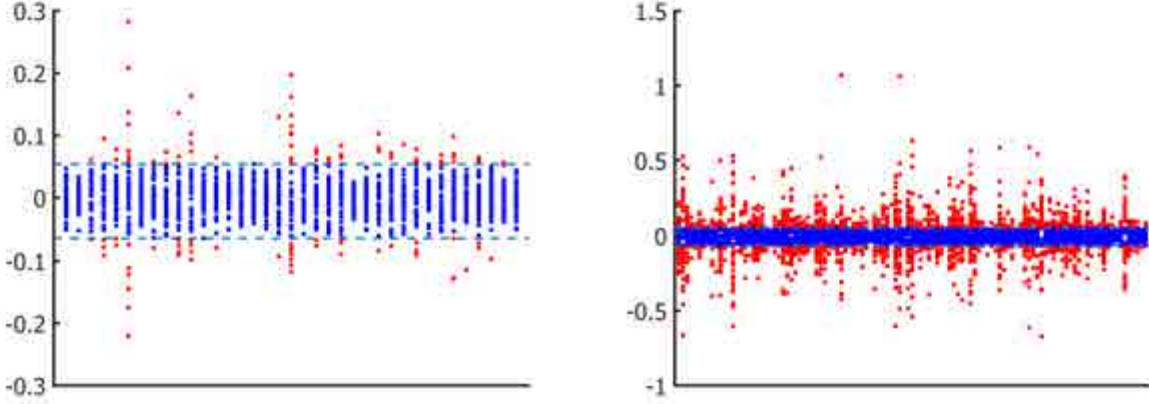$$\hat{\eta}_{it} = \pi_{i1}\hat{\eta}_{it-1} + \cdots \pi_{ip_i}\hat{\eta}_{it-p_i} + e_{it}. \tag{C.3}$$

Figure C.2: **Demeaned Residuals** $\left(\hat{e}_{it} - \overline{\hat{e}_i}\right)$**.** Outliers exceeding the $(-0.0685, 0.0778)$ band marked in red. Left: Energy intensity-Reducing Group. Right: Energy intensity-reducing group.

with the autoregressive order p_i determined by BIC for each country i. Denote parameter estimates and fitted residuals by $\hat{\pi}_{i1}, ..., \hat{\pi}_{ip_i}$ and $\hat{e}_{it}$ respectively. Re-sample from $(\hat{e}_{it} - \overline{\hat{e}_{i\cdot}})$ to obtain bootstrap samples (e_{it}^*) with outliers omitted as described in the Section C.2 below, from which bootstrap samples (η_{it}^*) are obtained using

$$\eta_{it}^* = \hat{\pi}_1^i \eta_{i,t-1}^* + \cdots \hat{\pi}_k^i \eta_{i,t-p_i}^* + e_{it}^*, \tag{C.4}$$

where η_{it}^* is set to $\hat{\eta}_{it}$ for $t = 1, ..., p_i$.

5. From equations (C.2) and (C.4), the bootstrapped values of the regressand are given by $Y_{it}^* = \hat{Y}_{it} + \hat{\lambda}_i' g_t^* + \eta_{it}^*$ with

$$\hat{Y}_{it} = \hat{\Gamma}_0(t, X_{it})X_{it} + \sum_{i=1}^{N} \hat{\alpha}_i \left[1 - \hat{\Gamma}_{1i}(t, X_{it})\right] + \sum_{i=1}^{N} \hat{\beta}_i \left[1 - \hat{\Gamma}_{2i}(t, X_{it})\right] X_{it}, \tag{C.5}$$

from the regression in (8) using the estimates obtained from the procedure described in Section 2.2.

6. Parameters and FPC loadings are re-estimated for each of 1,000 bootstrap replications of Y_{it}^* holding X_{it} and the FPCs fixed. The distributions of these parameter estimates are used to construct standard errors in Tables 1 and 2, and the distributions of the FPC loadings are used to construct confidence intervals for the loadings in Figures 4 and 5.
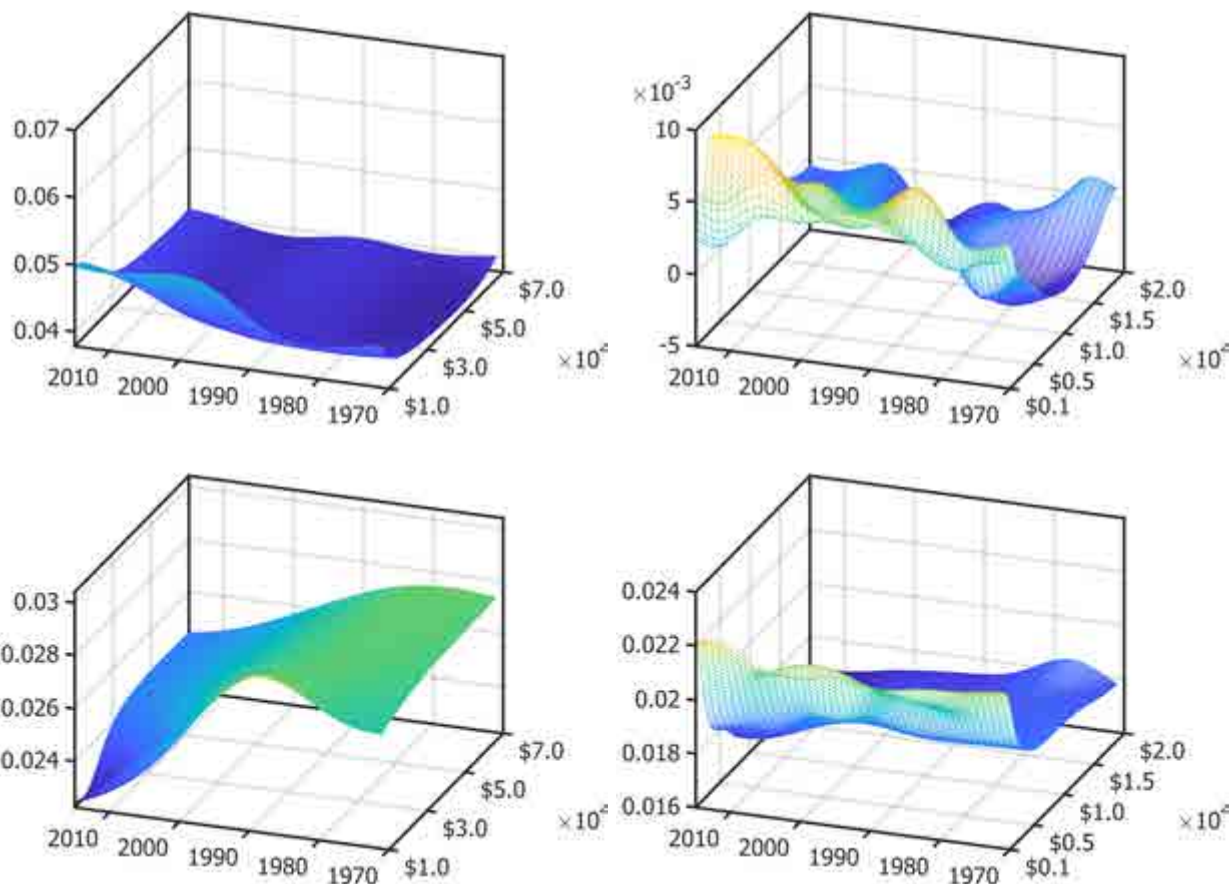
Figure C.3: **Bias and Standard Error Surfaces.** Top: Bias. Bottom: Standard errors. Left: Energy intensity-reducing group. Right: Energy intensity-reducing group.

C.2 Implementation

A few important details remain. 37 out of 46 countries in the intensity-reducing group and 113 out of 134 countries in the intensity-increasing group have data available over the full time span, 1971-2015. The regression in (C.1) is estimated for each of the two groups using these countries only for the purpose of the bootstrap. The eigenvalue ratio test of Ahn and Horenstein (2013) selects four factors common to the residuals of the 37 countries, which explain 81.3% of the variation of these residuals, and two factors common to the residuals of the 113 countries, which explains 56.6% of the variation in these residuals. For both groups, BIC selects $k = 1$ and $p_i \in [1, 5]$ with $p_i = 1$ most often.

In redrawing from $(\hat{e}_{it} - \overline{\hat{e}_i})$ for the 37 countries in the intensity-reducing group, we omit 64 outliers out of 1614 residuals. We similarly omit 1193 outliers out of 4939 residuals for the 113

countries in the intensity-increasing group.[10] Our aim in omitting outliers is to control excessive heterogeneity across countries resulting from major economic disruptions, such as armed conflict and the dissolution of the Soviet bloc and the Soviet Union, so the band $(-0.0685, 0.0778)$ is chosen based on the residuals of the G7 countries. The outliers are shown in Figure C.2.

The procedure just described omits countries with missing data. However, we use linear interpolation to fill the missing data for Kuwait over 1992-1994 following the Gulf War and for Niger in 1979. There are eight former Soviet or Soviet bloc countries in the intensity-reducing group for which we have incomplete time series, and there are 21 countries, some of which are former Soviet or Soviet bloc countries, for which that is the case in the intensity-increasing group.

We project the fitted residuals from the omitted countries in each group using common components obtained using the included countries in the respective group. The projection and its residuals may be labeled as $(\hat{\lambda}'_i \hat{g}_t)$ and $(\hat{\eta}_{it})$, and then we implement steps 3-5 of the bootstrap procedure using these projections and residuals.

For each group, we compute the bias and standard error of $\hat{\Gamma}(t, x_{it})$ from re-estimating $\Gamma(t, x_{it})$ over $1,000$ bootstrapped samples using both the countries with full data sets and those with missing data as described above. The bias is calculated as the difference between the mean of these $1,000$ estimates and the original estimate $\hat{\Gamma}(t, x_{it})$, while the standard error is calculated as the standard deviation of these $1,000$ estimates. The biases and standard errors for the two groups are plotted in Figure C.3. We interpret a positive bias of the bootstrap mean relative to the original estimate as a reflection of the magnitude of bias inherent in the estimation procedure, so that the bias-corrected surfaces shown in Figures 2 and 3 are the original surfaces $\hat{\Gamma}(t, x_{it})$ *minus* the biases shown in the top panels of Figure C.3.

---

[10]Note that $37 \times 45 = 1665$ so that $1665 - 1614 = 51$ for the energy intensity-reducing group and $113 \times 45 = 5085$ so that $5085 - 4939 = 146$ for the energy intensity-increasing group are equal to $\sum_{i=1}^{37} p_i$ and $\sum_{i=1}^{113} p_i$ respectively. These are the degrees of freedom used up by the lag orders.