A Neural Network Approach to the Environmental Kuznets Curve

Mikkel Bennedsen, Eric Hillebrand and Sebastian Jensen

CREATES Research Paper 2022-09

A Neural Network Approach to the Environmental Kuznets Curve

Mikkel Bennedsen[*], Eric Hillebrand[†], Sebastian Jensen[‡]

Abstract

We investigate the relationship between per capita gross domestic product and per capita carbon dioxide emissions using national-level panel data for the period 1960-2018. We propose a novel semiparametric panel data methodology that combines country and time fixed effects with a nonparametric neural network regression component. Globally and for the regions OECD and Asia, we find evidence of an inverse U-shaped relationship, often referred to as an environmental Kuznets curve (EKC). For OECD, the EKC-shape disappears when using consumption-based emissions data, suggesting the EKC-shape observed for OECD is driven by emissions exports. For Asia, the EKC-shape becomes even more pronounced when using consumption-based emissions data and exhibits an earlier turning point.

Keywords: Territorial carbon dioxide emissions; Consumption-based carbon dioxide emissions; Environmental Kuznets curve; Climate econometrics; Panel data; Machine learning; Neural networks.

JEL classification: C14, C23, C45, C51, C52, C53.

[*]Department of Economics and Business Economics and CREATES, Aarhus University, Fuglesangs Allé 4, 8210 Aarhus V, Denmark.

[†]Department of Economics and Business Economics and CREATES, Aarhus University, Fuglesangs Allé 4, 8210 Aarhus V, Denmark.

[‡]Department of Economics and Business Economics and CREATES, Aarhus University, Fuglesangs Allé 4, 8210 Aarhus V, Denmark. E-mail: smjensen@econ.au.dk

1   Introduction

Since the pre-industrial era, anthropogenic activity has driven atmospheric concentrations of green-house gases to reach levels unprecedented at least in the last 800,000 years (IPCC, 2014), and effects have been detected throughout the climate system. Carbon dioxide (CO_2) is the most important greenhouse gas and the key driver of climate change. It is estimated that CO_2 emissions are responsible for around 66% of total human contribution to temperature changes since 1750.[1] The observed increase in CO_2 emissions is largely driven by income and population growth (IPCC, 2014), highlighting the importance of a sound understanding of the relationship between per capita income and emissions. A better understanding of the relationship can provide important information to the public and policy makers alike, as it determines whether economic growth poses a threat to the environment, or rather presents a tool that can be used to combat climate change.

The aim of this paper is to investigate the shape of the income-emissions relationship, and to assess whether it varies across regions of the world or changes over time. We are particularly interested in possible evidence of an inverse U-shaped relationship, often referred to as an environmental Kuznets curve (EKC). This particular shape has played a prominent role in the literature and is surrounded with controversy. Using national-level panel data for the period 1960-2018, we estimate the relationship between per capita gross domestic product (GDP) and per capita CO_2 emissions. The model employed should be agnostic and able to balance cross-country dependencies and cross-country heterogeneity in the shape of the income-emissions relationship. We propose a novel semiparametric panel data methodology that combines parametric fixed effects with a non-parametric regression component consisting of a feedforward neural network. Some parameters of the neural network component are shared across all countries and used to learn common input transformations, while other parameters are specific to regions of homogeneous countries to allow for cross-region heterogeneity in the shape of the income-emissions relationship.

We propose a modeling framework that contains two distinct model specifications: a *static model* and a *dynamic model*. The static model specification contains both country and time fixed effects in addition to a neural network component that uses income as its only input variable. The model is static in the sense that the shape of the income-emissions relationship is assumed to be fixed over time. In this model, the level of the income-emissions relationship may change over time through the time fixed effects. The dynamic model, by contrast, uses a time variable as an additional input into the neural network component. By doing so, the dynamic model learns how the income-emissions relationship potentially changes its entire shape over time. One advantage of the dynamic model is also that it is easier to interpret the output of this model than that of a model with time fixed effects. We demonstrate, in a Monte Carlo experiment, that our proposed methodology is able to identify various functional forms of different complexity, and we also demonstrate its ability to account for country-specific stochastic trends. In addition, we demonstrate that the dynamic model is able to capture time-varying income-emissions relationships that cannot be captured using time fixed effects. When applying our methodology to empirical,

---

[1]Percentage is calculated based on Table 2 of NOAA (2020) and reflects concentration-based, anthropogenic radiative forcing for 2019 relative to 1750s. Radiative forcing measures the amount by which the Earth's energy budget is out of balance due to human activities.

territorial emissions data, we find evidence of an EKC-relationship globally and for the regions OECD and Asia. The dynamic model reveals that the global relationship is rather stable over time. On the other hand, it seems the EKC-shape does not appear until late in the sample period for OECD and Asia.

We also study the importance of international trade patterns for explaining the observed income-emissions relationships. We use two distinct types of CO_2 emissions estimates: territorial emissions (based on production) and consumption-based emissions. The latter accounts for cross-country emissions transfers through international trade. Comparing the relationships observed for these distinct types of emissions estimates allows us to asses the importance of international trade patterns. Using consumption-based emissions in our analysis instead of territorial emissions, the evidence of an EKC-relationship for OECD disappears. For Asia, the EKC-relationship becomes even more pronounced and with an earlier turning point. These findings suggest that the EKC-relationship observed for OECD when using territorial emissions is driven by emissions exports to other countries. The EKC-relationship observed for Asia when using territorial emissions is driven by local consumption and accounting for emissions imports results in an earlier turning point of the EKC.

The structure of this paper is as follows. In Section 2, we present the fundamental idea of the EKC and discuss econometric issues involved with its estimation. Section 3 summarizes the data set used in this paper. In Section 4, we present our proposed, neural network-based panel data methodology for analyzing the EKC for CO_2 emissions. Section 5 demonstrates the finite sample properties of our proposed methodology through a Monte Carlo experiment, and Section 6 presents the empirical analysis. Section 7 briefly summarizes and concludes.

2   The Environmental Kuznets Curve and its Econometric Issues

The literature on reduced-form modeling of the relationship between income and air pollution dates back to the seminal contribution by Grossman and Krueger (1991), who find evidence of an inverse U-shaped relationship between income and various air pollutants. The inverse U-shaped relationship between income and air pollution, sketched in Figure 1, is often referred to as the environmental Kuznets curve (EKC). Kuznets (1955) postulates a similar relationship between income and inequality. The EKC is consistent with the idea that, in early stages of economic development, air pollution grows rapidly as production expands (scale effect); however, as income grows, air pollution eventually reaches a turning point after which structural changes (composition effect) and technological improvements (technique effect) lead to a decline. The composition and technique effects may be driven by underlying factors such as increased environmental concerns and environmental policies, which in turn may be driven by even more fundamental factors. Empirically, support of the EKC hypothesis has been mixed. For instance, Selden and Song (1994), Schmalensee et al. (1998), Millimet et al. (2003), Costantini and Martini (2006), and Dutt (2009) find evidence in favor of the EKC hypothesis,[2] while, for example, Shafik (1994), Arrow et al. (1995), Holtz-Eakin

---

[2]We understand evidence of the EKC hypothesis to be evidence of an in-sample turning point in the estimated relationship between income and air pollution.

Figure 1: Illustrative environmental Kuznets curve
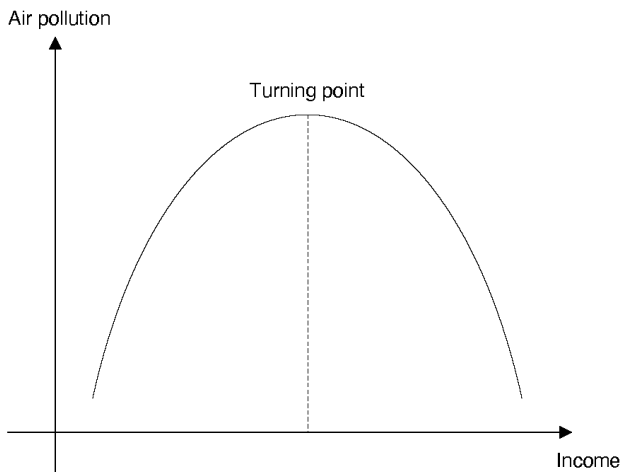
and Selden (1995), Wagner (2008), and Stern (2010) find evidence against it. For a recent survey of the EKC literature, see Stern (2017).

Since Grossman and Krueger (1991), the most commonly adopted modeling approach within the EKC literature has been to specify a panel data model consisting of country and time fixed effects in addition to some ad hoc polynomial specification (typically quadratic) of the functional relationship between income and air pollutant. However, this approach likely suffers from a variety of econometric issues (Stern, 2004; Müller-Fürstenberger and Wagner, 2007; Wagner, 2008; Aslanidis, 2009). Of these issues, the most critical seem to pertain to functional misspecification, cross-sectional heterogeneity, structural changes, non-identifiable time effects, integrated variables, and omitted variable bias. In this paper, we mainly focus on the issues of functional misspecification, cross-sectional heterogeneity, structural changes, and non-identifiable time effects, but also briefly address the issues of integrated variables and omitted variable bias.

Reliance on ad hoc parametric specifications of the functional relationship between income and air pollutant is problematic, as different parametric specifications can lead to significantly different conclusions, and functional misspecification is likely to occur (Harbaugh et al., 2002; Galeotti et al., 2006; Tsurumi and Managi, 2015). A strand of the EKC literature has therefore focused on semiparametric panel data models that combine use of parametric fixed effects with a nonparametric regression component. This strand has focused mainly on spline-based approaches (Schmalensee et al., 1998; Harbaugh et al., 2002; Dijkgraaf and Vollebergh, 2005; Auffhammer and Steinhauser, 2012) and kernel-based approaches (Taskin and Zaim, 2000; Millimet et al., 2003; Bertinelli and Strobl, 2005; Azomahou et al., 2006). Our work contributes to this discussion by suggesting a semiparametric panel data methodology where the nonparametric component consists of a feedforward neural network, one of the most popular machine learning algorithms. To the best of our knowledge, this is the first time machine learning methods have been applied to study the income-emissions relationship.

Assuming the true functional relationship between income and air pollution is of the EKC type,

the idea behind relying on national-level panel data is to use information from developing countries to provide information about the positively-sloped part of the EKC, and use information from developed countries to provide information about the negatively-sloped part of the EKC. However, some studies have challenged the assumption that all countries of the world would adhere to a single, global EKC, and advocate allowing for cross-country heterogeneity (de Bruyn et al., 1998; List and Gallet, 1999; Dijkgraaf and Vollebergh, 2005). Some studies also advocate the need to take into account structural changes over time (Halkos and Tsionas, 2001; Romero-Ávila, 2008; Hendry, 2018). We consider two distinct model specifications, distinguished by whether or not they can accommodate changes in the shape of the income-emissions relationship over time. We consider a static model specification, consisting of country and time fixed effects in addition to a neural network component that uses income as its only input variable. The static model imposes the assumption that the shape of the income-emissions relationship is fixed over time, although it does allow for level-shifts over time through the time fixed effects. We also consider a dynamic model specification. Instead of using time fixed effects to only allow for level-shifts in the income-emission relationship over time, it uses a time variable as an additional input into the neural network component. The dynamic model is able to learn how the income-emissions relationship potentially changes its entire shape over time. In both model specifications, we consider allowing for cross-sectional heterogeneity by letting a subset of parameters vary across regions of the world.

In the EKC literature, country and time fixed effects are usually treated as nuisance parameters, motivating either the random effects or the fixed effects transformation to remove them. However, such an approach might be problematic. The transformations remove only stochastic trends from the data that are common to all countries (Stern and Common, 2001). They cannot remove country-specific stochastic trends. In the presence of country-specific stochastic trends, methods relying on the random effects or the fixed effects transformation will be inconsistent. This situation is likely relevant, since country-specific stochastic trends are often encountered in practice (Perman and Stern, 1999). The importance of being able to properly account for stochastic trends is also emphasized by, for example, Perman and Stern (2003), Wagner (2008, 2015), and Lee and Lee (2009). Instead of treating country and time fixed effects as nuisance parameters and trying to remove them, we estimate fixed effects explicitly using a dummy variable approach. An extensive Monte Carlo experiment indicates that our proposed methodology is able to account for country-specific stochastic trends.

There is some evidence that the traditional EKC model omits important explanatory variables even after controlling for country and time-specific effects (Stern and Common, 2001; Magnani, 2001). To alleviate the problem of omitted variables, it has been suggested to extend the traditional EKC model with additional explanatory variables reflecting factors such as policies, governance, and institutions (Panayotou, 1997; Torras and Boyce, 1998; Dasgupta et al., 2006); education, urbanization, and inequality (Hill and Magnani, 2002; Barros et al., 2002; Borghesi, 2006); energy prices and fuel mix (Agras and Chapman, 1999; Richmond and Kaufmann, 2006), and international trade patterns (Suri and Chapman, 1998; Cole, 2003, 2004; Kearsley and Riddel, 2010). But, no canonical variables have been agreed upon in the literature thus far, and it is outside the scope of this paper to try to identify which variables to include to alleviate problems of omitted

4

variable bias. However, since international trade patterns are often considered one of the most important factors that can potentially explain the EKC (Dinda, 2004), we exploit the fact that we have available both territorial and consumption-based CO_2 emissions estimates, and compare the income-emissions relationship obtained using these distinct types of emissions estimates to asses the importance of international trade patterns. We are particularly interested in whether we find evidence of the much debated pollution haven hypothesis (PHH; Stern et al., 1996; Dinda, 2004): if changes in the structure of production experienced by developed countries are not met by similar changes in the structure of consumption, the EKC might just reflect displacement of emissions from developed countries with strong environmental regulation to less developed countries with weaker environmental regulation.

In summary, we propose a novel neural network-based panel data methodology for analyzing the EKC. To balance cross-country dependencies and cross-country heterogeneity in the shape of the income-emissions relationship, some parameters of the neural network are shared across all countries while other parameters are specific to regions of homogeneous countries. Our dynamic model specification presents a fundamentally new approach to dealing with the issue of non-identifiable separate income and time effects in a reduced form panel data model of the income-emissions relationship by using a time variable as an additional input into the neural network, instead of assuming separable income and time effects.

3   Data

To model the relationship between income and, respectively, territorial CO_2 emissions and consumption-based CO_2 emissions, we consider distinct, unbalanced panels of data. The specifications of the panels are reported in Table 1.

Territorial CO_2 emissions estimates are from the Global Carbon Project (2019). Estimates include emissions from fossil fuel combustion, oxidation, and cement production, and exclude emissions from bunker fuels, as the latter cannot be allocated unambiguously to particular countries. For the year 2018 (and for some countries 2016-2018), estimates are preliminary and made by the Global Carbon Project (GCP) based on energy statistics published by British Petroleum. For the first sample year 1960, the CO_2 panel covers 86 countries and accounts for 79.0% of the world's total population and 65.9% of the world's total CO_2 emissions. For the period 1990 onward, it covers more than 160 countries each year and accounts for more than 95.0% of the world's total population and CO_2 emissions.

Use of consumption-based CO_2 emissions in climate analysis has been advocated by Peters and Hertwich (2008a,b). To account for emissions transfers via international trade, consumption-based CO_2 emissions adjust territorial CO_2 emissions by adding emissions embedded in imports and subtracting emissions embedded in exports (Peters et al., 2011). Consumption-based CO_2 emissions estimates are also from the Global Carbon Project (2019). Throughout the sample period, which does not begin until 1990, the CO_2^C panel accounts for more than 90.3% of the world's total population and 93.5% of the world's total CO_2 emissions.

Table 1: Panel specifications

CO_2 panel	$CO_2{}^C$ panel	$CO_2{}^\star$ panel
• CO_2 emissions, Mt CO_2	• $CO_2{}^C$ emissions, Mt CO_2	• CO_2 emissions, Mt CO_2
• GDP, billion 2005 USD (PPP)	• GDP, billion 2005 USD (PPP)	• GDP, billion 2005 USD (PPP)
• Population, millions	• Population, millions	• Population, millions
• 1960-2018	• 1990-2017	• 1990-2017
• 186 (81) countries	• 117 (106) countries	• 117 (106) countries
• 8,641 observations	• 3,232 observations	• 3,232 observations

Note: The CO_2 panel contains data on territorial CO_2 emissions (CO_2 emissions) and is used in the territorial emissions analysis of Section 6.1; the $CO_2{}^C$ panel contains data on consumption-based CO_2 emissions ($CO_2{}^C$ emissions) and is used in the consumption-based emissions analysis of Section 6.2; and the $CO_2{}^\star$ panel is based on data from the CO_2 panel, restricted to the same sample period as the $CO_2{}^C$ panel, and is used as reference in the consumption-based emissions analysis of Section 6.2. Emissions are measured in megatonnes (Mt; i.e. 10^6 tonnes). Each panel also contains data on GDP measured in United States dollars (USD), adjusted using purchasing power parities (PPP), and population sizes measured in millions. Row four of the table reports sample lengths for each panel. Row five reports the number of countries in each panel, where numbers inside parentheses denote the number of countries with complete data, i.e. countries without "missing data" entries. The final row reports the total number of observations in each panel.

Data on population[3] and GDP are from the World Development Indicators database of the World Bank.[4] The GDP series that we consider is constructed from the separate series GDP[5] (current local currency units), the GDP deflator,[6] and a purchasing power parity (PPP) conversion factor.[7]

Besides estimating a global income-emissions relationship, we consider allowing for region-specific relationships. We specify regions in accordance with the macro-regions defined for the Shared Socioeconomic Pathways (SSPs) and related integrated assessment scenarios of the future (Riahi et al., 2017), used for example by the Coupled Model Intercomparison Project Phase 6 (CMIP6) of the World Climate Research Programme. The regions are defined as follows:

- OECD: 41 OECD90 and EU member states and candidates.

---

[3] The series "SP.POP.TOTL" was downloaded on December 12, 2019.

[4] Accessible at databank.worldbank.org/source/world-development-indicators, last accessed on December 12, 2019.

[5] The series "NY.GDP.MKTP.CN" was downloaded on December 12, 2019.

[6] The series "NY.GDP.DEFL.ZS" was downloaded on December 12, 2019.

[7] The series "PA.NUS.PPP" was downloaded on December 12, 2019.

- REF: 13 reforming economies of Eastern Europe and the former Soviet Union.

- Asia: 35 Asian countries excluding the Middle East, Japan and former Soviet Union states.

- MAF: 63 countries of the Middle East and Africa.

- LAM: 34 countries of Latin America and the Caribbean.

Table A.1 of the appendix maps each country to one of the five regions. Tables A.2 and A.3 of the appendix report descriptive statistics for the five regions. Table A.2 contains mean values for the data set. By some margin, OECD has been the richest region throughout the sample period. By contrast, Asia and MAF have been the poorest. In general, per capita emissions are the highest for the richest regions. Table A.3 contains standard deviations for the data set. OECD is the most heterogeneous region, followed by MAF. These two regions are also composed of the highest number of individual countries.

Table A.2 also indicates whether a region is net exporter or net importer of CO_2 emissions: if the level of consumption-based emissions is higher than for territorial emissions, it suggests the region is a net exporter of CO_2 emissions (and vice versa). By relying only on observations for which we have available both consumption-based and territorial emissions (compare means for $CO_2^\star$ and for CO_2^C in Table A.2), we aim to control for compositional differences in the CO_2 panel and the CO_2^C panel. Throughout the sample period, OECD is a net exporter of CO_2 emissions, while the regions REF, MAF, and Asia are net importers of CO_2 emissions. LAM goes from having a very balanced import and export of CO_2 emissions in the beginning of the sample period, to being a net importer of CO_2 emissions, to finally being a net exporter of CO_2 emissions. In the empirical analysis in Section 6, we discuss whether being a net exporter or net importer of CO_2 emissions can help explain some of the observed income-emissions relationships.

4   Methodology

The reduced-form relationship between income and greenhouse gas emissions can be mathematically represented in a panel data framework as

$$y_{it} = f\left(x_{it}, i, t\right) + u_{it}, \quad i = 1, \ldots, N_t,\ t = 1, \ldots, T, \tag{4.1}$$

where y_{it} is a measure of emissions, x_{it} is a measure of income, and u_{it} is an error term. Here, i indexes countries (or some other cross-sectional units) and t indexes time periods. Equation (4.1) allows for an unbalanced panel of data by letting $N_t \leq N$ denote the number of countries observed in time period t of the total number of countries N. The function f characterizes the functional relationship between income and emissions and is the main object of interest to us. Without imposing further restrictions, the function f cannot be identified from equation (4.1) as only one observation pair (y_{it}, x_{it}) is available for each combination of country and time period (i, t). Unfortunately, it is not clear which identifying restrictions are appropriate. As noted by Vollebergh et al. (2009), different ex-ante restrictions might be driving the non-robustness that plagues the literature on reduced-form estimation of the income-emissions relationship using panel

data, causing mixed empirical evidence of the EKC hypothesis. As discussed by Vollebergh et al. (2009), a particularly important but often overlooked issue relates to the fact that income and emissions are both time dependent. As both variables depend on time, separation of the effect of income from the effect of time hinges crucially on the restrictions used for identification. On the one hand, if one allows fully flexible time effects that are also cross-sectional specific, all variation in the data will be explained by these control variables. On the other hand, if one overly restricts time and cross-sectional effects, too much of the observed variation will be attributed to the income effect.

In the EKC literature, it is common to identify the function f in (4.1) by imposing the restriction that it is quadratic and common to all cross-sectional units and time periods, $f(x_{it}, i, t) = f(x_{it})$, and that u_{it} is a composite error term consisting of additively separable cross-section and time effects in addition to an idiosyncratic and purely stochastic effect. In this case, equation (4.1) reduces to a model of the form

$$y_{it} = \alpha_i + \beta_t + \delta_1 x_{it} + \delta_2 x_{it}^2 + \nu_{it}, \qquad \text{(Traditional EKC model)}$$

where α_i and β_t are treated either as random or fixed effects, and ν_{it} is a remainder stochastic error term. For now, we can think of ν_{it} as having a constant variance. Later, we loosen this assumption in Section 4.5. If α_i and β_t are correlated with income, a random effects model cannot be estimated consistently (Mundlak, 1978). Therefore, a fixed effects treatment is often preferred. The time and cross-section effects are supposed to control for omitted variables that are not endogenous consequences of income changes. The country fixed effects α_i are supposed to capture exogenous and persistent cross-sectional differences in features such as fossil fuel availabilities and prices, output mixes, regulatory structures, policies, and tastes. The time fixed effects β_t are supposed to capture time-varying omitted variables and shocks that are common to all countries (Stern, 2017). They capture effects on emissions over time in absence of changes in income, and we mainly interpret them as effects of common technology shocks that are not captured by income changes. If omitted variables are correlated with income and not properly controlled for by the country and time fixed effects, then the slope coefficients δ_1 and δ_2 will capture both direct and indirect effects of changes in income on emissions. In a typical EKC study, y_{it} is log-transformed per capita emissions of a greenhouse gas like CO_2 and x_{it} is log-transformed per capita GDP. The focus on per capita quantities reflects the hypothesis that population sizes do not affect average behavior. The log-linear specification is typically preferred over a linear specification as multiplicative cross-sectional and time specific effects are deemed more plausible than additive effects given the heterogeneity of cross-sectional units in a typical study (Schmalensee et al., 1998). Linear specifications have also been considered in the literature, but no large differences between the two specifications are typically observed (Holtz-Eakin and Selden, 1995).

Although the traditional EKC model has been extensively applied within the literature, the appropriateness of it has been heavily debated (Stern, 2004, 2017; Müller-Fürstenberger and Wagner, 2007; Wagner, 2008). Rather than pre-imposing the restrictions that f is necessarily quadratic and common to all cross-sectional units and time periods, and that u_{it} is a composite error term consisting of additively separable cross-section, time, and idiosyncratic effects, we consider a number of different identifying restrictions and their implications for the estimated income-emissions

relationship. Throughout, we drop the parametric restriction that f is necessarily quadratic. Our goal is to represent f using a feedforward neural network, imposing as few parametric restrictions as possible on equation (4.1).

In the next sections, we present our proposed methodology. In Section 4.1, we propose a static neural network model of the form

$$y_{it} = \alpha_i + \beta_t(r) + f^{\text{NN}}(x_{it}, r) + \nu_{it}, \qquad \text{(Static neural network model)}$$

where the superscript on f highlights that we use a neural network to model f nonparametrically, and r is a regional indicator. A function $\mathfrak{r} : \{1, \ldots, N\} \to \{1, \ldots, R\}$, $1 \le R \le N$, is initially used to map each country to a macro-region within which countries are assumed to admit the same time effects and functional relationship f (up to country-specific intercept shifts). This model specification retains the restriction that u_{it} is a composite error term consisting of additively separable country and time effects in addition to an idiosyncratic effect, but drops the restriction that f is necessarily quadratic and common to all countries. The term *static* used in the name of this model specification refers to the fact that the unknown function of interest f is assumed to be time-invariant in line with the traditional approach of the literature.

In Section 4.2, we propose a dynamic neural network model of the form

$$y_{it} = \alpha_i + f^{\text{NN}}(x_{it}, t, r) + \nu_{it}, \qquad \text{(Dynamic neural network model)}$$

which drops the restriction of additively separable time effects. Instead, f is allowed to depend explicitly on t. In this way, we leave it up to the neural network to learn how income and time together affect emissions, and the entire shape of f is allowed to change over time. This is the sense in which this model specification is *dynamic*.

4.1 Static neural network model

The static neural network model imposes the restrictions that the shape of the functional relationship between income and emissions f is time-invariant, and that the error term u_{it} is a composite error term consisting of additively separable country and time fixed effects in addition to an idiosyncratic and purely stochastic effect. However, we impose no parametric restrictions on the functional form of f. Instead, we learn f directly from the data using a feedforward neural network. We allow for cross-sectional heterogeneity by initially mapping each country to a region of comparable countries according to $\mathfrak{r} : \{1, \ldots, N\} \to \{1, \ldots, R\}$, $1 \le R \le N$, then allow a subset of model parameters to vary across regions. The model can be mathematically represented as

$$y_{it} = \alpha_i + \beta_t(r) + \phi(r)^{\top} z_{it}^{(h)} + \nu_{it}, \quad i = 1, \ldots, N_t, \ t = 1, \ldots, T, \ r = \mathfrak{r}(i), \qquad (4.2)$$
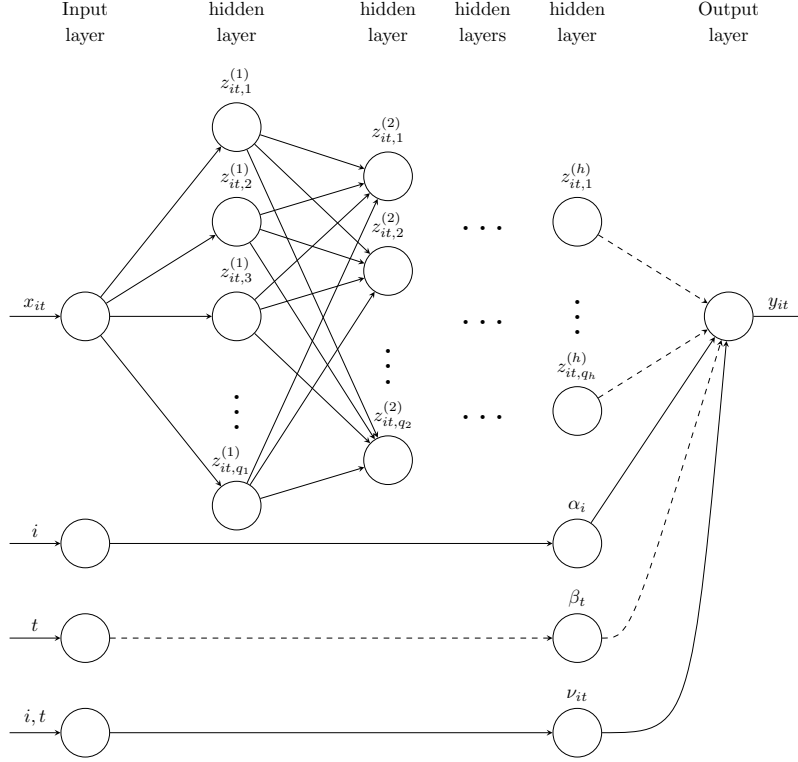
$$z_{it}^{(h)} = g\left(\kappa^{(h)} + \Gamma^{(h)} z_{it}^{(h-1)}\right), \qquad (4.3)$$

$$z_{it}^{(h-1)} = g\left(\kappa^{(h-1)} + \Gamma^{(h-1)} z_{it}^{(h-2)}\right), \qquad (4.4)$$

$$\vdots$$

$$z_{it}^{(1)} = g\left(\kappa^{(1)} + \Gamma^{(1)} x_{it}\right), \qquad (4.5)$$

9

Figure 2: Static neural network model

Note: The notation $z_{it,j}^{(\ell)}$ refers to the j-th element of the vector of derived variables in the ℓ-th layer for country i at time t. Dashed edges in the the graph represent region-specific connections.

where $x_{it}, y_{it} \in \mathbb{R}$ are the natural logarithm of, respectively, per capita GDP and per capita emissions of country i in year t. The vector $\phi(r) \in \mathbb{R}^{q_h}$ is a column vector of region-specific slope coefficients, and $z_{it}^{(\ell)} \in \mathbb{R}^{q_\ell}$, $\ell = 1, \ldots, h$, is a column vector of variables derived by the neural network. Equation (4.2), which we will refer to as the output layer of the model, is a linear regression model in derived variables $z_{it}^{(h)}$ augmented by country fixed effects α_i and region-specific time fixed effects $\beta_t(r)$. The derived variables constitute nonlinear transformations of the input variable x_{it}, learned through so-called hidden layers of the neural network given by (4.3)–(4.5). Instead of y_{it} depending linearly on x_{it} and x_{it}^2 as in the traditional EKC model, we learn optimal input transformations directly from the data. Note that parameters of the hidden layers are common to all countries. The neural network exploits cross-country dependencies by using the entire panel of data to learn optimal input transformations. Nevertheless, we let slope coefficients of the output layer $\phi(r)$ be region-specific to allow the shape of the income-emissions relationship to vary across regions. Different choices for the number of regions R provide different bias-variance tradeoffs by relying on different degrees of parameter sharing across countries. At the one extreme, $R = 1$ results in a global model with one, global income-emissions relationship (shape) and only level shifts across countries (country fixed effects). At the other extreme, $R = N$ results in a national model where each country has its own income-emissions relationship. In between those extremes,

$1 < R < N$, we have a regional model with R region-specific income-emissions relationships. In the empirical analysis of Section 6, we consider $R = 1$, $R = N$, and $R = 5$, where the five regions considered are OECD, REF, Asia, MAF, and LAM, defined in Section 3. In equations (4.2) and (4.3), $h \geq 1$ denotes the number of hidden layers used to learn the derived variables and is said to determine the *depth* of the neural network; q_ℓ denotes dimensionality of the ℓ-th hidden layer and is referred to as *width*. This implies that $z_{it}^{(\ell)}$ is a q_ℓ-dimensional column vector of derived variables, $\kappa^{(\ell)}$ is a q_ℓ-dimensional column vector of unknown intercepts[8] to be estimated from data, and $\Gamma^{(\ell)}$ is a $q_\ell \times q_{\ell-1}$ coefficient matrix[9] to be estimated from data. We follow the convention to set $q_0 \equiv 1$, implying that $\Gamma^{(1)}$ is a q_1-dimensional column vector. The non-linear function g is referred to as an *activation function*. The nonlinearity of g is what allows the network to learn nonlinear relationships. We rely on the Swish activation function of Ramachandran et al. (2017), which is found to outperform the standard rectified linear unit (ReLU) activation function (Glorot et al., 2011) on a number of different tasks (Ramachandran et al., 2017). The Swish function is defined as $g(z) = z \left(1 + \exp(-z)\right)^{-1}$ for $z \in \mathbb{R}$, and can be considered a smoothed version of the ReLU function defined as $g(z) = \max(z, 0)$. We prefer the Swish function as it provides a smoother f in the static model where only one input variable is passed through the hidden layers of the neural network.

The choices of overall network depth and width of each hidden layer constitute important architectural considerations with implications for the representation capabilities of the network. It can be shown that even a shallow feedforward network with just one hidden layer can approximate any Borel-measurable function from one finite-dimensional space to another to any desired degree of accuracy provided the network is wide enough (Hornik et al., 1989; Cybenko, 1989; Leshno et al., 1993). However, adding depth to a network allows it to learn increasingly more abstract representations of the input data, and, typically, less wide networks with fewer parameters overall are required for a given level of accuracy when compared to a shallow network. On the other hand, deeper networks are typically more difficult to optimize than shallow ones, and require more input data to avoid overfitting to in-sample noise (Goodfellow et al., 2016). In Section 4.4, we discuss how we decide on an optimal architecture.

The static model is visually represented in Figure 2 by means of a directed acyclic graph (DAG). The upper part of the graph illustrates the neural network component. The bottom part of the graph illustrates the parametric component. Edges represent how information travels through the model. The neural network is feedforward as information only travels forward through the network without any feeedback loops. Vertices of the input layer reflect information that is presented to the model. Vertices of the hidden layers reflect elements of the sequentially derived z-variables. Remaining vertices represent elements of the error term and the output.

4.2 Dynamic neural network model

For the dynamic neural network model, we drop the restriction of a time-invariant functional relationship f, and instead let the neural network learn how the functional relationship depends

---

[8]Typically referred to as *biases* in the neural network literature.

[9]Typically referred to as *weights* in the neural network literature.

on time. We do so by dropping the assumption of time fixed effects and instead pass a time variable through the hidden layers of the neural network together with income. Again, we allow for cross-sectional heterogeneity by initially mapping each country to a region of comparable countries according to $\mathfrak{r} : \{1, \dots, N\} \to \{1, \dots, R\}$, $1 \leq R \leq N$, then allow a subset of model parameters to vary across regions. The model can be mathematically represented as

$$y_{it} = \alpha_i + \phi(r)^\top z_{it}^{(h)} + \nu_{it}, \quad i = 1, \dots, N_t, \ t = 1, \dots, T, \ r = \mathfrak{r}(i), \tag{4.6}$$

$$z_{it}^{(h)} = g\left(\kappa^{(h)} + \Gamma^{(h)} z_{it}^{(h-1)}\right), \tag{4.7}$$

$$z_{it}^{(h-1)} = g\left(\kappa^{(h-1)} + \Gamma^{(h-1)} z_{it}^{(h-2)}\right), \tag{4.8}$$

$$\vdots$$

$$z_{it}^{(1)} = g\left(\kappa^{(1)} + \Gamma^{(1,1)} x_{it} + \Gamma^{(1,2)} t\right), \tag{4.9}$$

with $x_{it}, y_{it}, z_{it}^{(\ell)}, \phi, \kappa^{(\ell)}, \Gamma^{(\ell)}$ as above, and $\Gamma^{(1,1)}$, $\Gamma^{(1,2)}$ being q_ℓ-dimensional column vectors of unknown slope coefficients. Once again, (4.6) is a linear regression model in derived variables $z_{it}^{(h)}$, learned through hidden layers of a neural network, augmented by country fixed effects α_i. What is new is that $z_{it}^{(h)}$ is derived from both income and time. In the static model formulation, $z_{it}^{(h)}$ is derived from income only. We note that, because there is only one observation pair (x_{it}, y_{it}) available for each country at each point in time, we cannot sensibly consider a national formulation

Figure 3: Dynamic neural network model
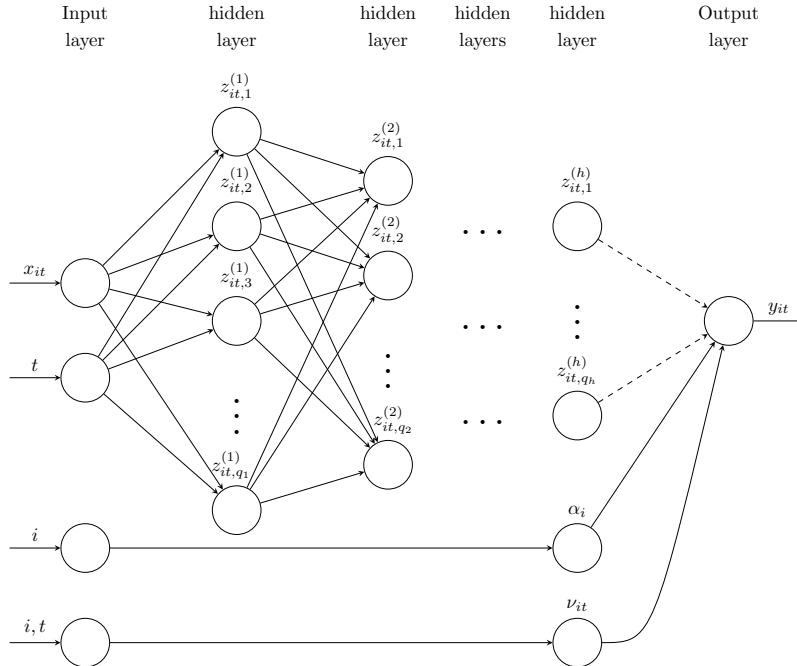
Note: The notation $z_{it,j}^{(\ell)}$ refers to the j-th element of the vector of derived variables in the ℓ-th layer for country i at time t. Dashed edges in the upper part of the graph represent region-specific connections.

$(R = N)$ of the dynamic model as this would result in a perfect fit. The DAG associated with the dynamic model is shown in Figure 3. By passing a time variable through the hidden layers of the neural network together with income, the network is able to learn how the two variables interact, and hence how the functional relationship between income and emissions potentially changes shape over time.

4.3   Estimation

For both the static and dynamic model, we estimate[10] all free parameters simultaneously by minimizing the sum of region-specific mean squared errors:

$$J(\phi, \kappa, \Gamma, \alpha, \beta) = \sum_{r=1}^{R} \sum_{t=1}^{T} \sum_{i \in I_r} \frac{1}{n_r} \left(y_{it} - \hat{y}_{it}(x_{it}, i, t) \right)^2, \tag{4.10}$$

where n_r is the number of observations for region r, and $I_r \subseteq \{1, 2, \ldots, N\}$ is the set of indices of countries belonging to region r. Note the model output $\hat{y}$ is dependent on the estimated parameter vectors $\hat{\phi}$, $\hat{\kappa}$, $\hat{\Gamma}$, $\hat{\alpha}$, and possibly $\hat{\beta}$, which are understood to contain model parameters of all relevant layers and regions. A possible alternative to the loss function in (4.10) would be to minimize global mean squared errors where each contribution to the loss function is weighted by $1/n$ instead of $1/n_r$. However, by specifying the loss function as a sum of region-specific mean squared errors, we incorporate a tradeoff between the fits on each region. We ensure the estimated model has an accurate fit for each region separately and that the model is not biased toward regions with more observations than others. Note that the weights have no effect on parameter estimates in the global $(R = 1)$ formulation of the models.

Before discussing practical details concerning minimization of J, we briefly elaborate on a subtlety regarding identification of fixed effects. The flexibility of the neural network causes it to automatically adapt its predictions to the level of the dependent variable in the estimation sample. From a purely illustrative point of view, we can think of the neural network as implicitly identifying an overall intercept even if not specified explicitly. This is true even if imposing restrictions such as $f(0) = 0$. For this reason, we can only identify country fixed effects for all but one country and time fixed effects for all but one time period. In all applications of this paper, we use the U.S. as a reference when estimating a global model formulation $(R = 1)$, then include dummies for the remaining countries. Hence, estimated income-emissions relationship will be normalized for the U.S., and estimated country fixed effects should be interpreted as intercept shifts relative to the U.S. When estimating a regional model formulation $(R = 5)$, we must account for the fact that we are in a sense estimating five neural networks simultaneously. Therefore, we must normalize the income-emissions relationships for each region separately. We use the U.S. as reference for OECD, Russia as reference for REF, China as reference for Asia, South Africa as reference for MAF, and Mexico as reference for LAM. When estimating a national formulation $(R = N)$ of the static model, we cannot identify any country fixed effects. When estimating the static model, we also use the initial time period $t = 1$ as reference, then include dummies for the remaining

---

[10]Often referred to as *training* in the neural networks literature.

time periods $t = 2, \ldots, T$. Hence, estimated time fixed effects should be interpreted as intercept shifts relative to the initial time period. When estimating a regional formulation ($R = 5$) of the static model, we allow region-specific time fixed effects. When estimating a national formulation ($R = N$) of the static model, we exclude time fixed effects: since there is only one observation pair (x_{it}, y_{it}) available for each country at each point in time, including (country-specific) time fixed effects would result in a perfect fit.

We follow the standard in the neural networks literature and minimize J using gradient descent. Since the optimization problem is high-dimensional and nonconvex, the loss function most likely features numerous critical points. Whereas most second-order methods would get stuck in any critical point, gradient descent is often able to escape critical points associated with a high loss (Goodfellow et al., 2016). However, naive gradient descent performs poorly whenever the Hessian matrix has a poor condition number. Without modification, gradient descent does not know to prefer directions of parameter space where the slope of the loss function remains negative for longer. Poor conditioning of the Hessian matrix also makes it difficult to determine a good step size for the gradient descent algorithm.[11] One must balance the goals of avoiding to overshoot the minimum and being able to make significant progress in directions with little curvature. We use the popular Adam (ADAptive Moment estimation) variant of gradient descent. It seeks to mitigate these issues through the use of momentum and a separate learning rate for each parameter that is automatically adapted at each iteration of the algorithm (Kingma and Ba, 2014). The Adam algorithm is illustrated in Table A.5 of the appendix. For each iteration[12] of the Adam algorithm, we use the entire data set when evaluating the gradient.[13] We use suggested defaults for the hyperparameters of the Adam algorithm (Kingma and Ba, 2014). We find they work well in all applications. We stop the Adam algorithm when we have observed no significant decrease in the loss function over 100 consecutive iterations, using a tolerance level of 10^{-6}, then restore parameter estimates associated with the lowest loss across all iterations (not necessarily the last iteration).

Since Adam is based on local moves, it works well only if initialized within a well-behaved region of parameter space that is connected to a satisfactory solution by a path the algorithm can follow. It is therefore important to use appropriate initial values for α, β, ϕ, κ, and γ. We initialize all slope coefficients randomly from a truncated normal distribution as suggested by He et al. (2015). This is the suggested way to break symmetry (Goodfellow et al., 2016), motivate each hidden unit to learn a different function, and avoid problems of vanishing and exploding gradients when using ReLU-type activation functions. We use multiple different initializations to ensure the minimization routine has arrived at a satisfactory minimum. All intercepts (including fixed effects) are initialized from zero as is standard in the neural networks literature (Goodfellow et al., 2016).

4.4 Model selection

The choice of neural network architecture is a standard model selection task: there is a tradeoff between choosing large values for the depth and width parameters to reduce bias, and choosing

---

[11]Often referred to as the *learning rate* in the neural networks literature.

[12]Often referred to as an *epoch* in the neural networks literature.

[13]Often referred to as *batch learning* in the neural networks literature.

small values to ensure smoothness. Within the neural networks literature, hypothesis tests, cross-validation, Bayesian regularization, dropout, and information criteria have been proposed for model selection (Anders and Korn, 1999). We use the Bayesian information criterion (BIC). In the Monte Carlo experiments of Section 5, we show the BIC identifies a close approximation to all functions we consider. The BIC takes the form

$$\text{BIC} = \log J(\hat{\phi}, \hat{\kappa}, \hat{\Gamma}, \hat{\alpha}, \hat{\beta}) + \frac{p \log n}{n}, \tag{4.11}$$

where $J(\hat{\phi}, \hat{\kappa}, \hat{\Gamma}, \hat{\alpha}, \hat{\beta})$ is the objective function in (4.10) evaluated in parameters estimates, n is total number of observations, and p is total number of model parameters excluding fixed effects.

As discussed by Gu et al. (2020), it is not necessary to search over unreasonably many network architectures. We fix a reasonable number of architectures of varying complexity ex ante. We focus on rectangular and pyramid-shaped architectures as proposed by Masters (1993), which are useful for learning gradually more abstract transformations of the input variable and for keeping the number of free parameters at a reasonable level. The full list of candidate architectures is presented in Table A.4 of the appendix together with the number of free model parameters.

4.5   Prediction intervals

Quantifying the uncertainty associated with neural network predictions is often neglected as it is generally a very challenging task. Zapranis and Livanis (2005) reviews attempts based on either bootstrapping and other ensemble techniques, analytical approaches based on nonlinear regression theory in a nonparametric setting, or maximum likelihood techniques. More recently, a promising strand of the literature has evolved that considers the use of approximate Bayesian inference techniques based on Markov Chain Monte Carlo (MCMC) or variational inference (e.g. Blundell et al., 2015). In this paper, we propose quantifying the uncertainty stemming from noise in the estimation data following the maximum likelihood approach of Nix and Weigend (1994b,a).

The methodology presented above can be understood as estimating the expected value of a conditional target distribution. One could also try to estimate higher-order information about the conditional target distribution. In the empirical analysis of Section 6, we assume the conditional target distribution is Gaussian, then seek to obtain an additional estimate of the variance, as suggested by Nix and Weigend (1994b,a). Specifically, we consider a model of the form

$$y_{it} = m(x_{it}, i, t) + s(x_{it}, i)\nu_{it}, \ \ \nu_{it} \sim \mathcal{N}(0, 1).$$

The function $m(x_{it}, i, t)$ captures the expected value of the conditional target distribution and is the function of inherent interest. The term $s(x_{it}, i)\nu_{it}$ captures noise around the expected value with standard deviation $s(x_{it}, i)$ that allows for heteroskedasticity. With estimates available for both the expected value and the standard deviation of the conditional target distribution, we construct (conditional) 95% pointwise prediction intervals (PIs) as $\hat{m}(x_{it}, i, t) \pm 1.96\hat{s}(x_{it}, i)$. These prediction intervals are conditional in the sense that they ignore model uncertainty.

In the empirical analysis of Section 6, we consider both a static model specification for $m(x_{it}, i, t)$,

$$m(x_{it}, i, t) = \alpha_i + \beta_t(r) + f_m^{\text{NN}}(x_{it}, r),$$

and a dynamic model specification for $m(x_{it}, i, t)$,

$$m(x_{it}, i, t) = \alpha_i + f_m^{\text{NN}}(x_{it}, t, r),$$

but retain a static model specification for $s(x_{it}, i)$ with fixed effects replaced by a region-specific intercept,

$$s(x_{it}, i) = \tilde{g}\left(\kappa(r) + f_s^{\text{NN}}(x_{it}, r)\right),$$

where $\tilde{g}$ is a numerically stable activation function used to ensure nonnegativity of $s(x_{it}, i)$, implemented using $\tilde{g}(z) = 10^{-6} + \log(\exp(0.001z) + 1)$ for $z \in \mathbb{R}$. The neural networks f_m^{NN} and f_s^{NN} are specified separately with no shared connectivity as suggested by Nix and Weigend (1994b). As discussed by Nix and Weigend (1994a), one could alternatively connect some of the hidden units across f_m^{NN} and f_s^{NN} if there is an expected overlap in the mappings from inputs to respectively the expected value and the standard deviation of the conditional target distribution. As presented above, we use the BIC to select the architecture of the network f_m^{NN}. For the network f_s^{NN}, we specify a simple one-hidden-layer architecture with two hidden units.

To estimate the parameters of $m(x_{it}, i, t)$ and $s(x_{it}, i)$, we use a two-step procedure as suggested by Zapranis and Livanis (2005). In the first step, the parameters of $m(x_{it}, i, t)$ are estimated by minimizing the cost function in (4.10). In the second step, the parameters of $s(x_{it}, i)$ are estimated by minimizing a similar cost function based on the squared residuals from the first step:

$$J_s(\phi_s, \kappa_s, \Gamma_s, \alpha_s, \beta_s) = \sum_{r=1}^{R} \sum_{t=1}^{T} \sum_{i \in I_r} \frac{1}{n_r} \left[(y_{it} - \hat{y}_{it}(x_{it}, i, t))^2 - \hat{s}^2(x_{it}, i)\right]^2,$$

where $\hat{y}_{it}(x_{it}, i, t) = \hat{m}(x_{it}, i, t)$ is estimated in the first step.

5 Monte Carlo Experiment

In this section, we demonstrate finite-sample properties of our proposed methodology in a controlled setup that closely resembles the situation in the empirical analysis of Section 6. We demonstrate that our proposed methodology is able to identify various functional forms and account for country-specific stochastic trends. We also demonstrate that it is important to use the dynamic model in situations where the true relationship of interest is time-varying.

5.1 Static neural network model

To investigate finite-sample properties of the static neural network model, we simulate data from the following data-generating process:

$$y_{it} = f(x_{it}, r) + u_{it}, \qquad i = 1, \dots, N_t, \ t = 1, \dots, T, \ r = \mathfrak{r}(i),$$

$$u_{it} = \alpha_i + \beta(r)\log t + \nu_{it}, \quad \nu_{it} \overset{iid}{\sim} \mathcal{N}(0, \sigma_\nu^2),$$

where $\mathfrak{r}: \{1, \dots, N\} \to \{1, \dots, R\}$ is a function mapping each country to one of the R regions. To obtain as realistic a setup as possible, where the input variables are characterized by multiple

16

Table 2: Optimal neural network architectures for the static model in the Monte Carlo experiment

	Function		
	Linear	Quadratic	Cubic
Global model formulation			
Neural network architecture	(2)	(4)	(4)
# parameters (excl. fixed effects)	4	12	12
Regional model formulation			
Neural network architecture	(2)	(4)	(4)
# parameters (excl. fixed effects)	14	28	28

Note: "(a)" indicates a neural network architecture with one hidden layer containing a units.

country-specific stochastic trends, x_{it} is taken to be the natural logarithm of per capita GDP from the CO_2 emissions panel discussed in Section 3. We consider a global model formulation ($R = 1$) and a regional model formulation with $R = 5$. In line with the discussion of Section 3, we consider the five regions OECD, REF, Asia, MAF, and LAM. We focus on these model formulations as they are the main focus in Section 6. We consider the following different specifications of f:

$$f(x_{it}, r) = \delta_1(r)x_{it}, \qquad\qquad\qquad \text{(Linear function)}$$
$$f(x_{it}, r) = \delta_1(r)x_{it} + \delta_2(r)x_{it}^2, \qquad\qquad\qquad \text{(Quadratic function)}$$
$$f(x_{it}, r) = \delta_1(r)x_{it} + \delta_2(r)x_{it}^2 + \delta_3(r)x_{it}^3, \qquad\qquad\qquad \text{(Cubic function)}$$

where $\delta_1(r), \delta_2(r), \delta_3(r) \in \mathbb{R}$ are region-specific, constant parameters. To obtain realistic parameter values, we use estimates from simple regressions based on the function under consideration (linear, quadratic or cubic) and the data discussed in Section 3. We use the natural logarithm of per capita CO_2 emissions from the CO_2 panel of Section 3 as dependent variable and the natural logarithm of per capita GDP from the CO_2 emissions panel and appropriate transformations (square and cube) as input variables. We include dummy variables for each country and a logarithmic time trend. Note that time fixed effects can be backed out from the estimated time trend. For the regional model formulation, we interact input variables and the time trend with a regional indicator to allow region-specific slope coefficients and time effects. Finally, we scale some estimated slope coefficients to achieve a desired amount of variation in f. We set $\sigma_\nu^2 \equiv 0.35$, which is the empirical standard deviation of the residuals from the simple regression using the natural logarithm of per capita GDP and its square as input variables, without any interactions. For simplicity, we use this value throughout all experiments in the section.

In order to demonstrate the finite sample properties of the proposed estimation procedure, we simulate 100 Monte Carlo samples for each model formulation, global and regional, and the three functions. To keep runtime at a reasonable level, we also construct an initial Monte Carlo sample that is used to determine the optimal network architecture. This architecture, chosen on

17

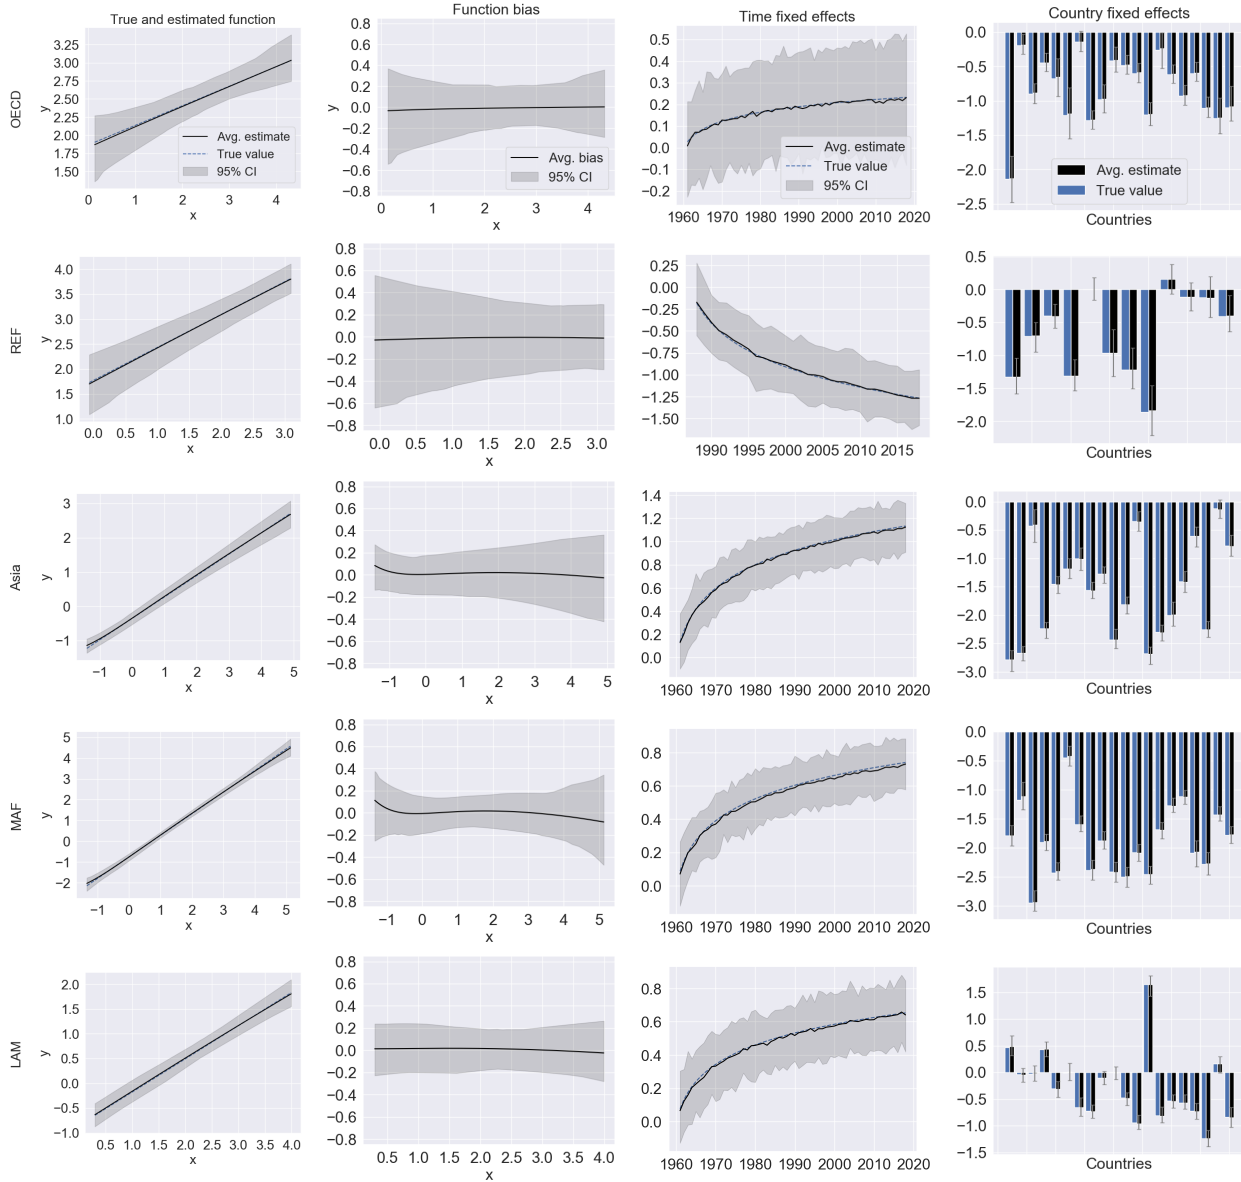Figure 4: Monte Carlo results for the global formulation of the static model
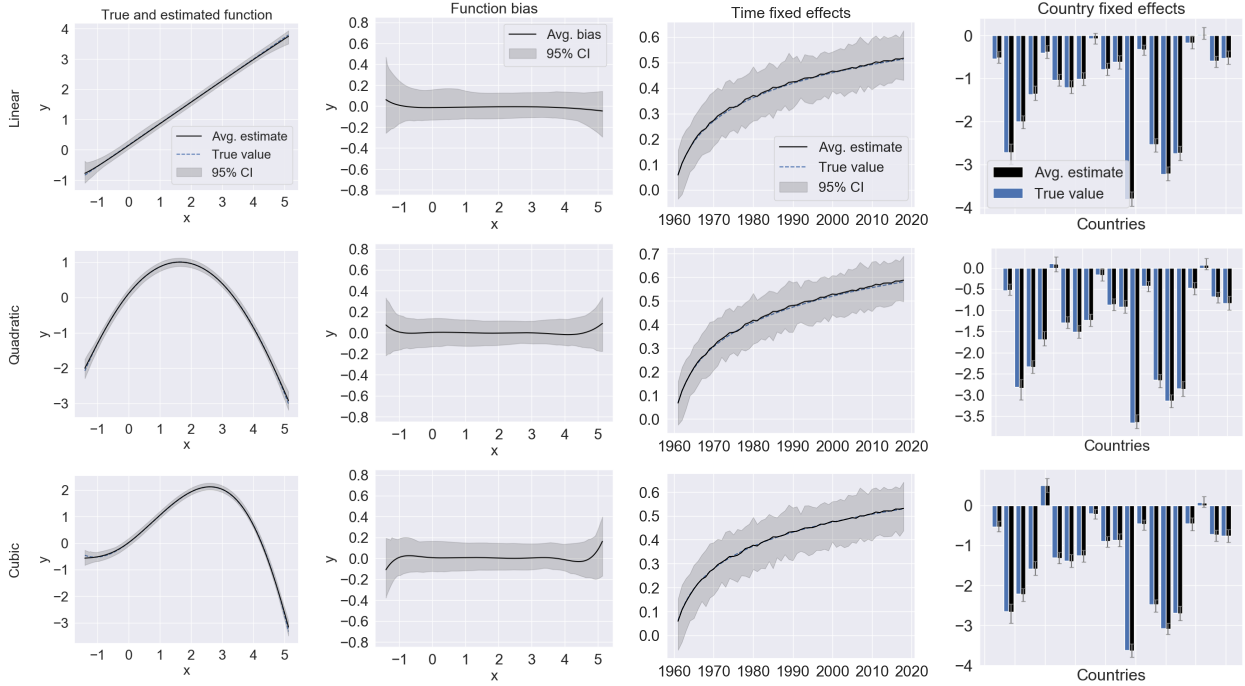
Note: In the plots in columns 1 and 3, the true values are covered by the average estimates throughout most of the input region. In the plots in column 4, gray error bars indicate 95% confidence intervals.

the initial sample, is then kept for the remaining 100 samples. In this way, the results presented here can be seen as a lower bound on the performance of the estimators, since their properties would likely improve if the optimal architecture was allowed to be chosen separately for each individual Monte Carlo sample. In Table 2, we report optimal network architectures for both the global and the regional model formulation, determined by minimizing the BIC (4.11) on an initial Monte Carlo sample. As discussed in the previous section, we search over the set of architectures presented in Table A.4 of the appendix. For both model formulations, global and regional, and all functions, the optimal neural network architecture is a simple one-hidden-layer architecture. For the linear function, two hidden units are optimal for both the global and regional model formulation. However, since the regional model formulation has region-specific slope coefficients in the output layer, it has more free parameters to estimate than the global model formulation for a given network architecture, even without considering the added number of time fixed effects associated with the regional formulation. For the quadratic and cubic function, four hidden units are optimal for both model formulations. Note that it is optimal to slightly increase the complexity of the network architecture, and so the approximation capabilities of the model, as we increase the complexity of the function to approximate. Note also that the models contain a large number of free parameters. For instance, we see from Table 2 that the regional formulation of the static model with only a single hidden layer containing four hidden units contains 28 free parameters excluding fixed effects.

For the global model formulation, average estimates across Monte Carlo samples are reported

Figure 5: Monte Carlo results for the regional formulation of the static model: linear function

Note: In the plots in columns 1 and 3, the true values are covered by the average estimates throughout most of the input region. In the plots in column 4, gray error bars indicate 95% confidence intervals.

in Figure 4. The first column of Figure 4 contains the true functional relationships and average estimate across Monte Carlo samples, plotted over the entire input region used for estimation. We also plot 95% confidence bands obtained from the quantiles of the estimates across Monte Carlo samples. We note that the model is able to capture all three functions to a high degree of accuracy with narrow confidence bands. In the second column of Figure 4, we plot function bias, i.e. the difference between the curves shown in the first column. We initially note the bias is around zero for all three functions throughout most of the input region. We also note the model displays largest discrepancies toward the boundaries of the input region used for estimation. This is likely related
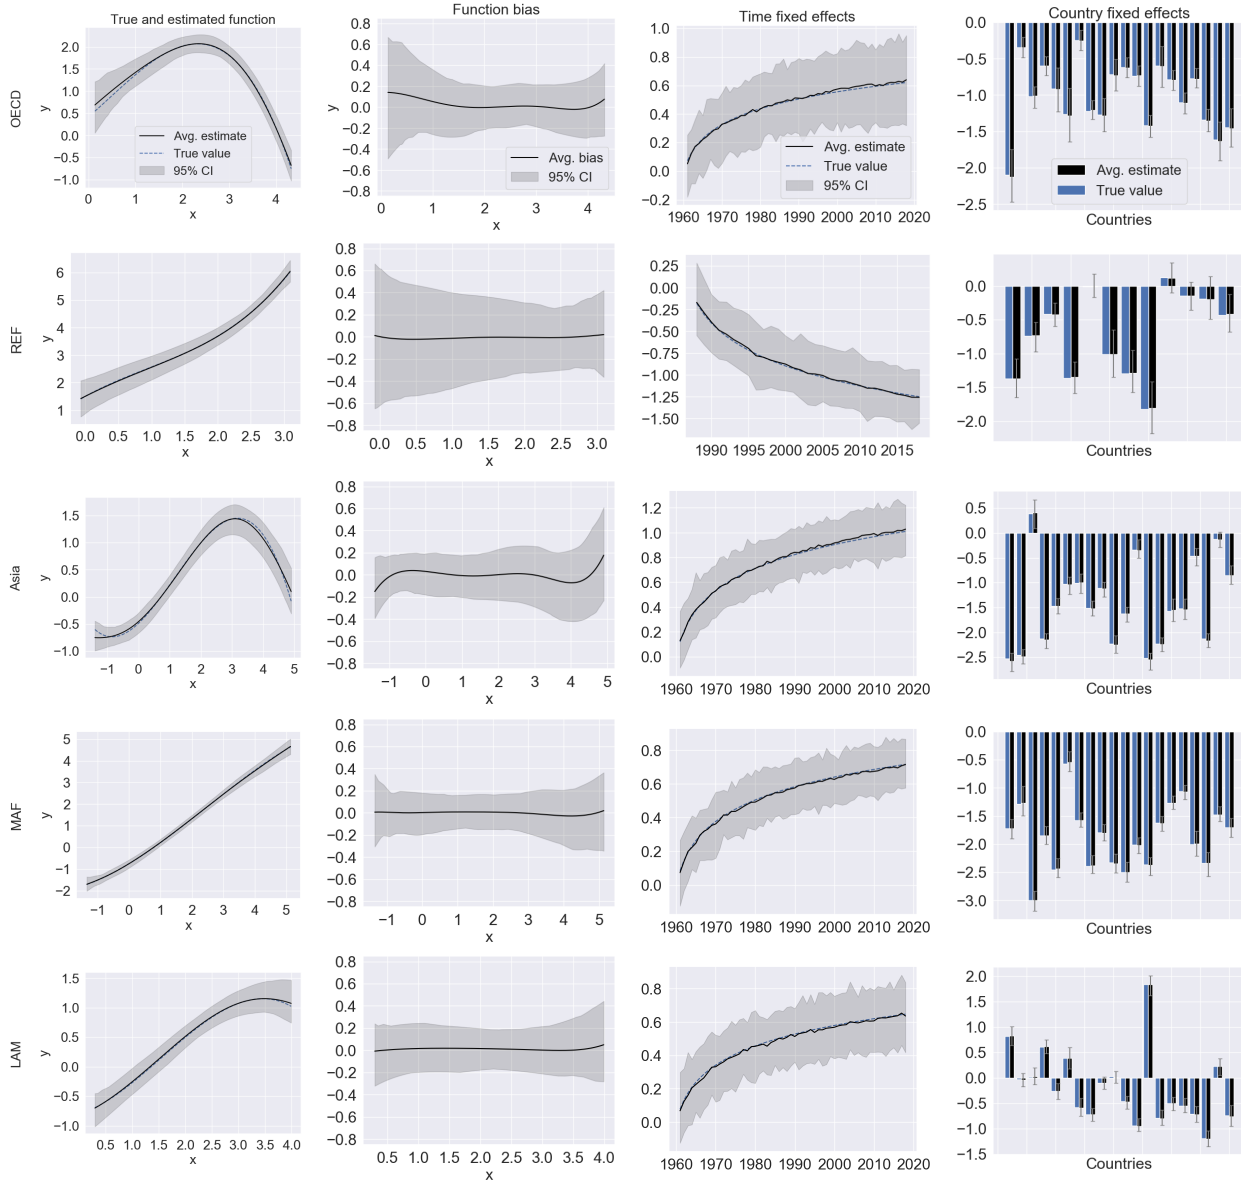
Figure 6: Monte Carlo results for the regional formulation of the static model: quadratic function
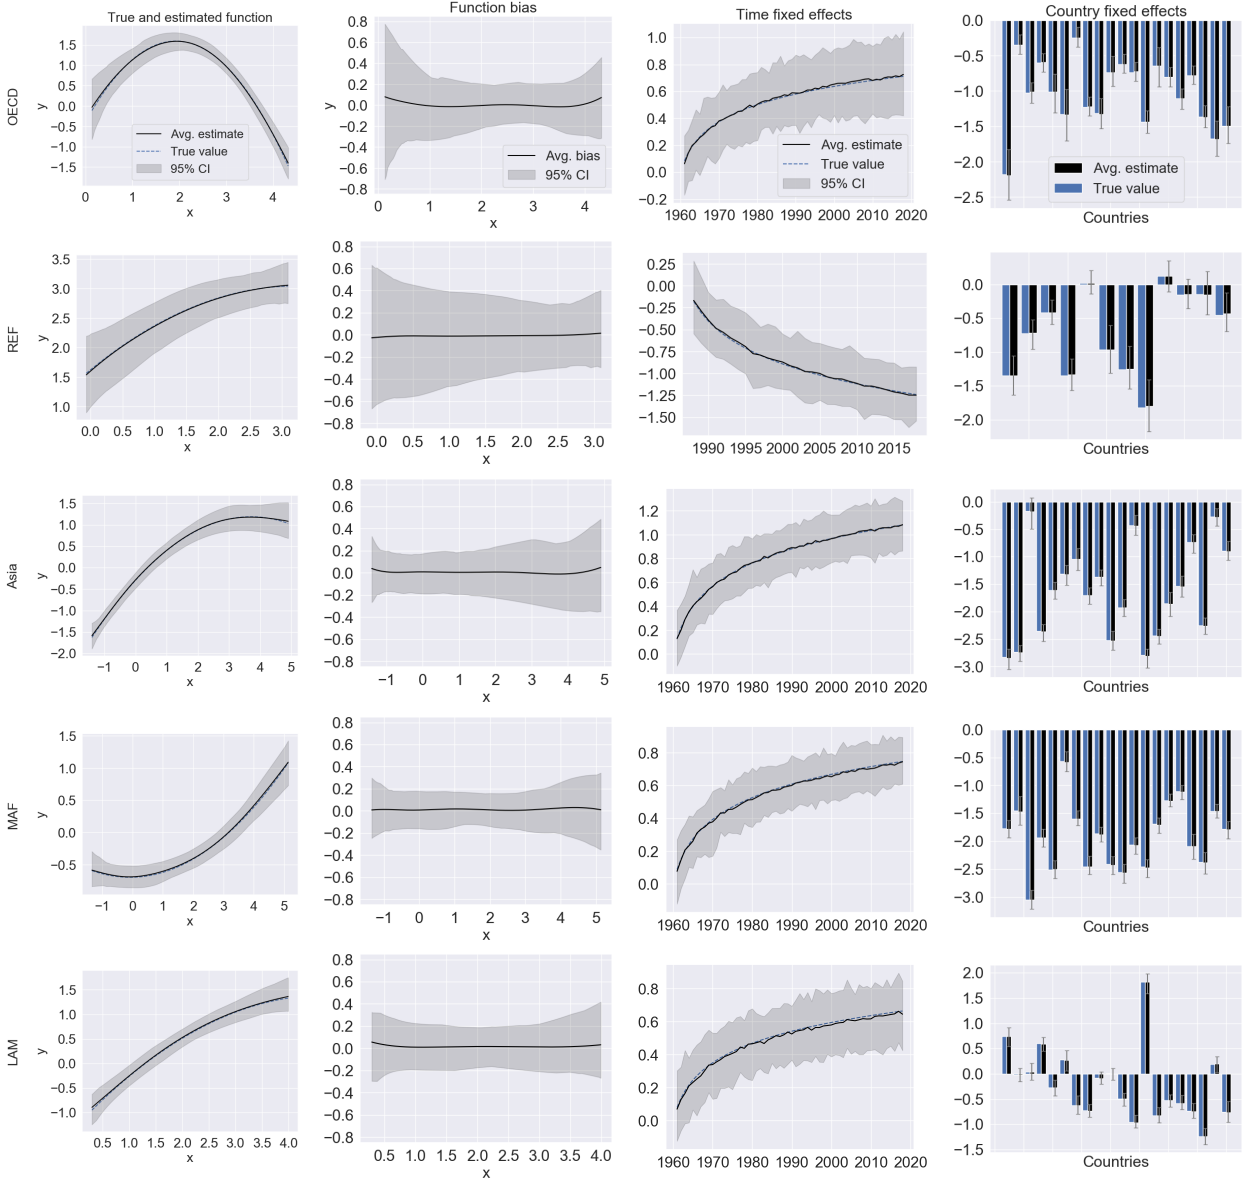
Note: In the plots in columns 1 and 3, the true values are covered by the average estimates throughout most of the input region. In the plots in column 4, gray error bars indicate 95% confidence intervals.

to the sparsity of observations toward the boundaries of the input region used for estimations and the so-called boundary issues well known from the literature on nonparametric regression using kernel-based techniques (Malec and Schienle, 2014). Although a correction mechanism is outside the scope of this paper, it suggests that one should be careful when interpreting model output based on input values close to or beyond the boundary of the input region used for estimation. The third and fourth columns of Figure 4 contain true values and average estimates for the time and country fixed effects, respectively. Again, we note the model is able to capture the true values with a high degree of accuracy for all three functions.

Figure 7: Monte Carlo results for the regional formulation of the static model: cubic function

Note: In the plots in columns 1 and 3, the true values are covered by the average estimates throughout most of the input region. In column 4, gray error bars indicate 95% confidence intervals.

Figures 5, 6, and 7 summarize average estimates across Monte Carlo samples for the regional model formulation and the linear, quadratic, and cubic function, respectively. Although the learning task faced by the regional model formulation seems more complex than that faced by the global model formulation, as the regional model formulation is faced with the task of learning region-specific functions and time effects, the accuracy of the regional model formulation seems on par with that of the global model formulation in Figure 4 for each region separately. It is also encouraging that the model seems to perform about equally well for all regions despite large degree of variation in the number of observations available for each region. For instance, the model has

available as little as 373 observation for the region REF and as many as 3008 observations for the region MAF.

5.2 Dynamic neural network model

To investigate finite sample properties of the dynamic neural network model, we simulate data from the following data-generating process:

$$y_{it} = f(x_{it}, t, r) + u_{it}, \quad i = 1, \dots, N_t, \ t = 1, \dots, T, \ r = \mathfrak{r}(i),$$

$$u_{it} = \alpha_i + \nu_{it}, \qquad \nu_{it} \overset{iid}{\sim} \mathcal{N}(0, \sigma_\nu^2),$$

with notation as above. We again consider $R = 1$ (global model) and $R = 5$ (regional model). In contrast to the static model, the effect of time is not modeled by a logarithmic time trend. Instead, f now changes its entire shape over time. Specifically, we consider the following different specifications of f:

$$f(x_{it}, t, r) = \delta_1(r)x_{it}, \qquad \text{(Linear function)}$$

$$f(x_{it}, t, r) = \frac{T - t + 1}{T}\delta_1(r)x_{it} + \frac{t - 1}{T}\delta_2(r)x_{it}^2, \qquad \text{(Linear-quadratic function)}$$

$$f(x_{it}, t, r) = \frac{T - t + 1}{T}\delta_1(r)x_{it} + \frac{t - 1}{T}\delta_2(r)x_{it}^2 + \frac{t - 1}{T}\delta_3(r)x_{it}^3, \qquad \text{(Linear-cubic function)}$$

where $\delta_1(r), \delta_2(r), \delta_3(r) \in \mathbb{R}$ are region-specific, constant parameters. To obtain realistic parameter values, we once again resort to estimates from simple regressions based on the function under consideration (linear, linear-quadratic or linear-cubic). Using the linear-cubic function as an example, we run a regression using the natural logarithm of per capita CO_2 emissions from the CO_2 panel of Section 3 as dependent variable and the natural logarithm of per capita GDP from the CO_2

Table 3: Optimal neural network architectures for the dynamic model in the Monte Carlo experiment

| | Function | | |
	Linear	Linear-quadratic	Linear-cubic
Global model formulation			
Neural network architecture	(2)	(4,4)	(4,2,2)
# parameters (excl. fixed effects)	8	36	30
Regional model formulation			
Neural network architecture	(2)	(4,4)	(4,4)
# parameters (excl. fixed effects)	16	52	52

Note: "(a,b,c)" indicates a neural network architecture with three hidden layers containing a units in the first layer, b in the second, and c in the third.

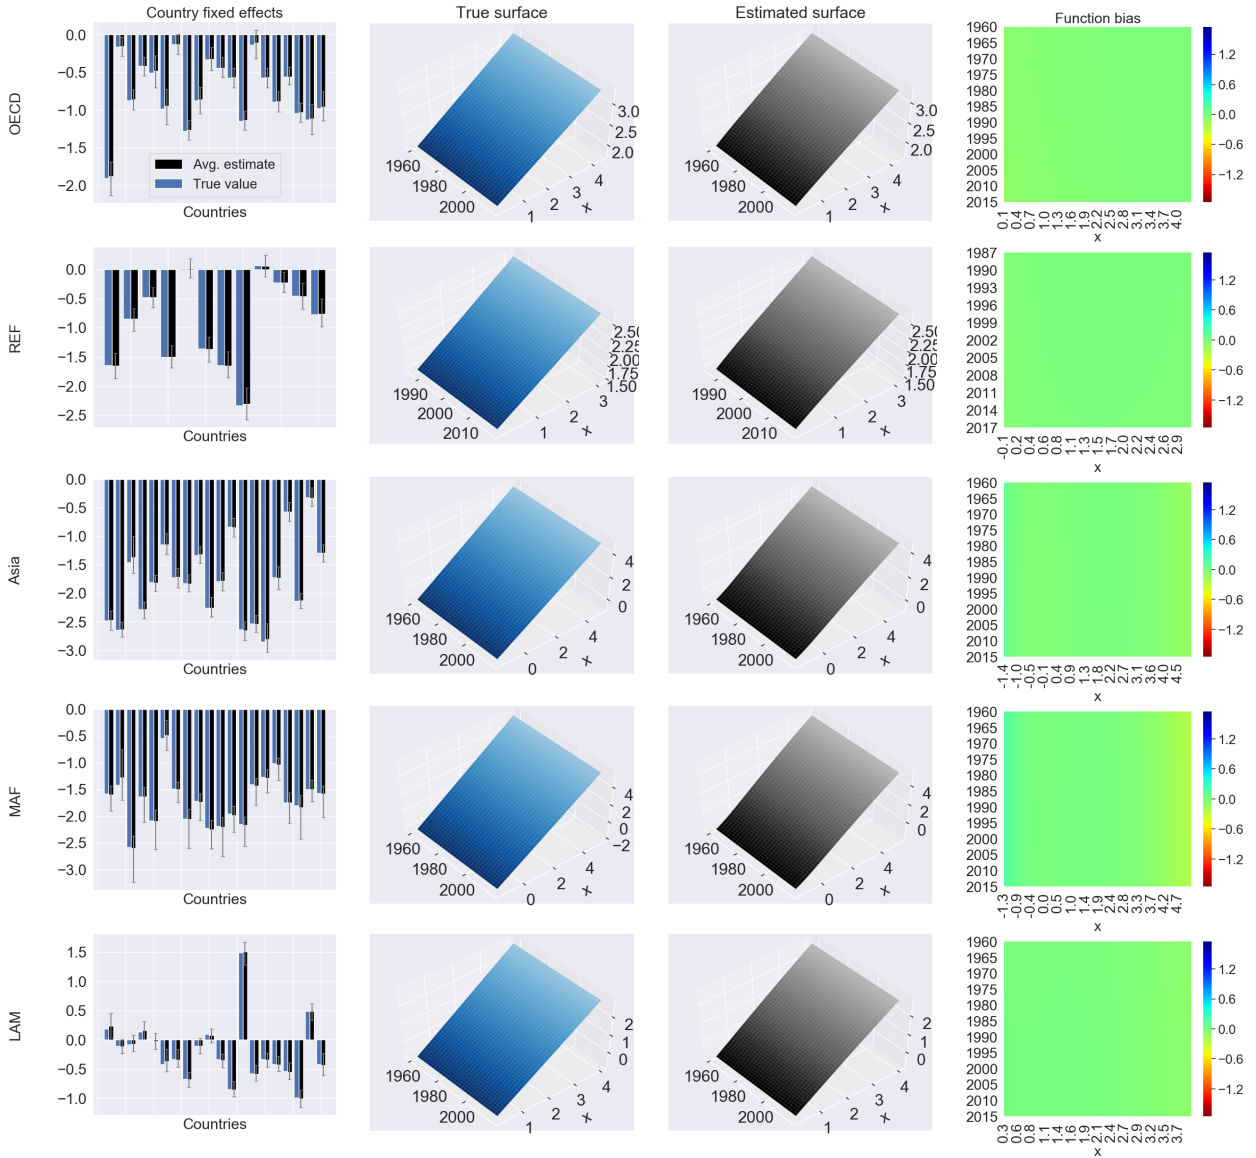Figure 8: Monte Carlo results for the global formulation of the dynamic model
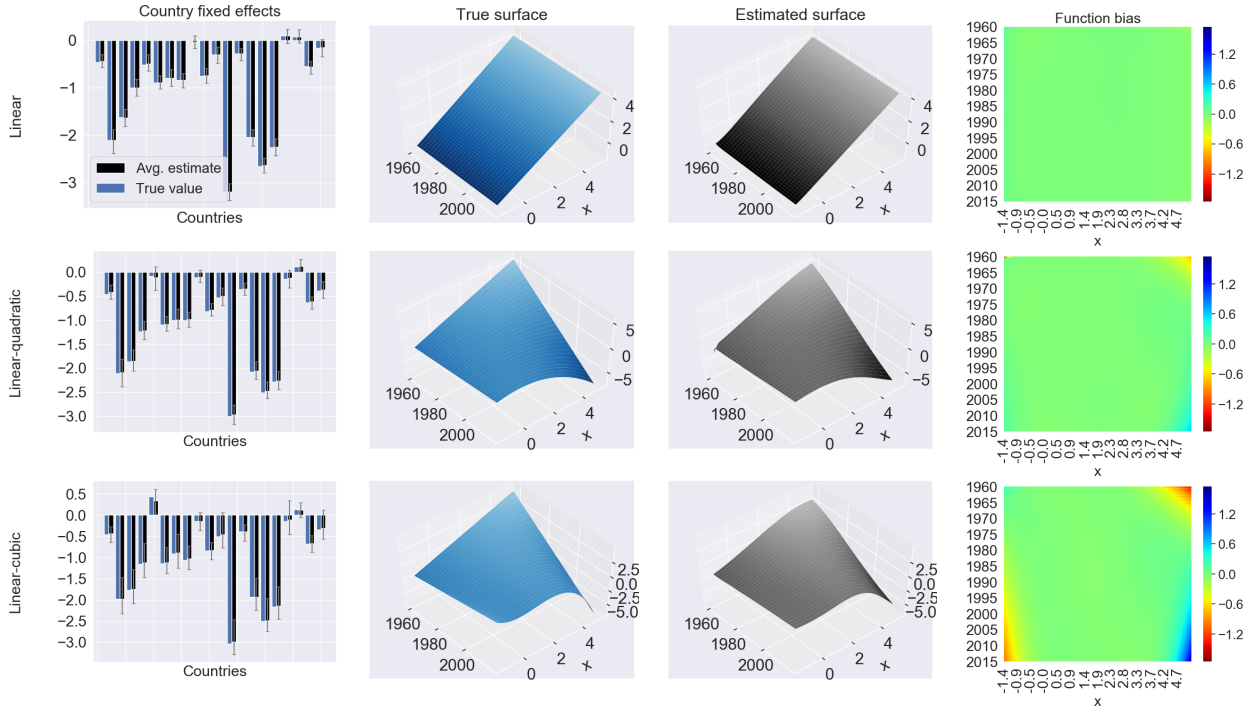
Note: In the country fixed effects plots, gray error bars indicate 95% confidence intervals.

emissions panel and its square and cube as input variables to obtain realistic values for the constant parameters $\delta_1(r), \delta_2(r)$ and $\delta_3(r)$. We scale some estimated slope coefficients to achieve a sufficient amount of variation in f. For the regional model formulation, we interact input variables with a regional indicator to allow region-specific slope coefficients. We also include dummy variables for each country in the regression to obtain realistic values for α_i.

Table 3 contains optimal network architectures for both the global and regional formulation of the dynamic model obtained by minimizing the BIC (4.11) on an initial Monte Carlo sample as discussed above. For the simple linear function, both model formulations require only a single-hidden-layer architecture with two hidden units, the simplest architecture in the set of candidates. Not counting the number of fixed effects, the dynamic model contains more free parameters than its static counterpart for a given formulation and network architecture because both time and income are passed through the hidden layers of the neural network. The learning task faced by the dynamic model seems in many ways more complex than that of its static counterpart. The dynamic model is required to learn how the functional relationship of interest changes its entire shape over time, whereas the static model allows only for intercept shifts over time. From Table 3, we see that increased complexity of the learning task prompts an increase in optimal model complexity. In particular, it is optimal to add depth to the network architecture when turning from the linear function to the more complex and time-varying linear-quadratic and linear-cubic functions.
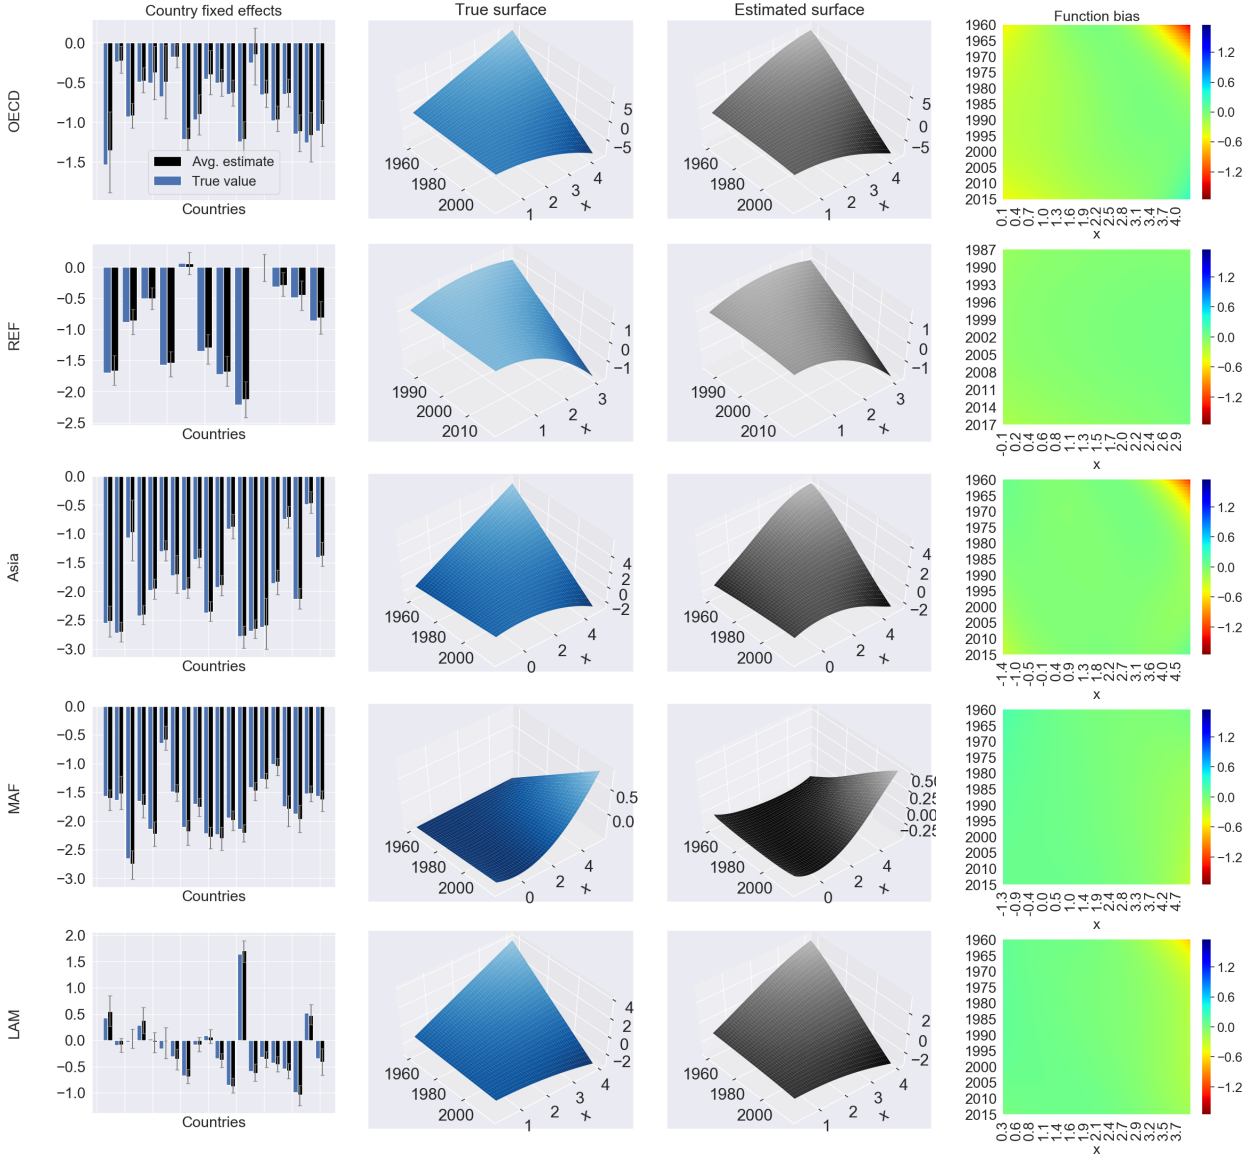
Figure 9: Monte Carlo results for the regional formulation of the dynamic model: linear function

Note: In the country fixed effects plots, gray error bars indicate 95% confidence intervals.

Figure 8 shows the finite sample properties of the global formulation of the dynamic model by summarizing average estimates across 100 Monte Carlo samples. The first column of plots shows true and estimated country fixed effects together with 95% confidence intervals obtained from the quantiles of the estimates across Monte Carlo samples. For all three functions, the model is able to accurately estimate the country fixed effects. The second and third columns of plots show, respectively, true and estimated surfaces, plotted over the entire input region used for estimation. The fourth column of plots displays heat maps of function bias obtained by subtracting the surfaces in the second and third columns. Initially, we note the bias is around zero throughout most of the input region used for estimation. However, similar to the static model, it is more difficult for the model to accurately estimate the true surface toward the boundaries of the input region.

24

Figure 10: Monte Carlo results for the regional formulation of the dynamic model: linear-quadratic function

Note: In the country fixed effects plots, gray error bars indicate 95% confidence intervals.

Indeed, the worst performance of the model occurs toward the four corners of the input region. The boundary issues seem most severe for the linear-cubic function, likely because most of the cubic curvature becomes the most pronounced toward the boundaries of the input region.

Figures 9–11 summarize average estimates across Monte Carlo samples for the regional formulation of the dynamic model for the linear, linear-quadratic, and linear-cubic function, respectively. Despite large variations in the number of observations available for each region, the model performs well for each region and for every function considered.

Figure 11: Monte Carlo results for the regional formulation of the dynamic model: linear-cubic function
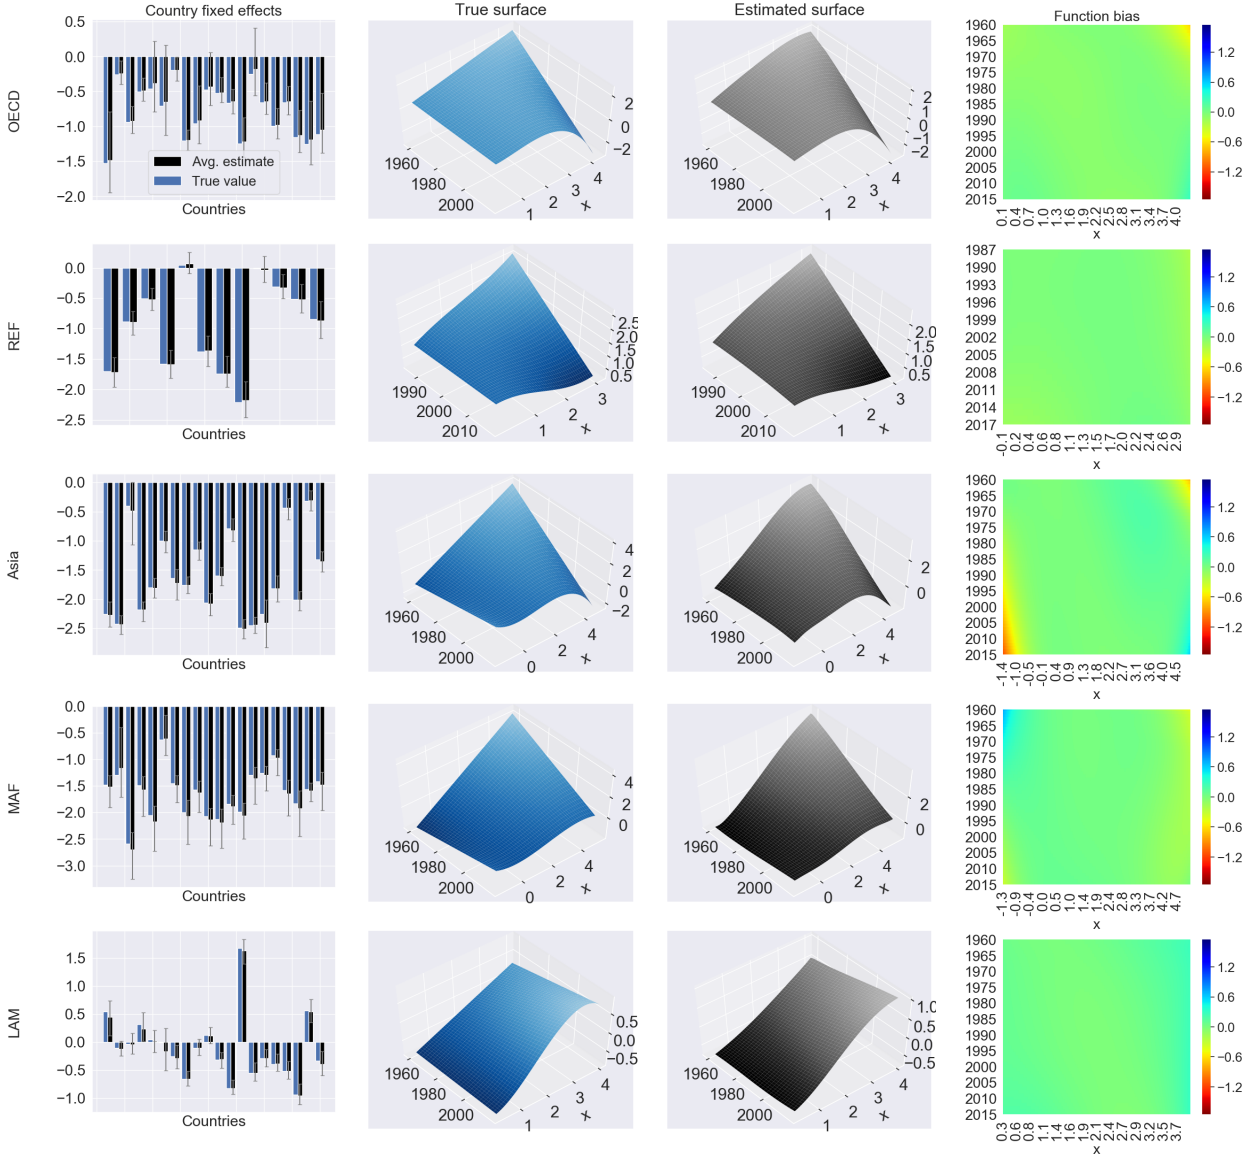
Note: In the country fixed effects plots, gray error bars indicate 95% confidence intervals.

5.3 Misspecification

In this section, we investigate effects of model misspecification on our estimator of the parameters in the static model. We are interested in the performance of the static model in cases where the true income-emissions relationship is time varying. We simulate observations using the "linear-quadratic" and "linear-cubic" dynamic specifications presented above, then estimate a global ($R = 1$) and regional ($R = 5$) formulation of the static model. The underlying simulations are thus the same as those used in Section 5.2, but we now apply the estimation method used in Section 5.1. For ease of presentation, we focus on results for the global model formulation. We report optimal

Figure 12: Monte Carlo misspecification results for the global formulation of the static model

Note: In the country fixed effects plots, gray error bars indicate 95% confidence intervals.

network architectures based on an initial Monte Carlo sample in Table A.6 of the appendix.

Note that while the static model estimates a one-dimensional income-emissions relationship, we can construct an implied, two-dimensional, estimated income-emissions surface by adding estimated time fixed effects to the estimated income-emissions function. Figure 12 shows estimation results for the global model formulation. Results for the regional model formulation are similar and are reported in Figures A.1 and A.2 of the appendix. The first column in Figure 12 shows true and estimated country fixed effects; the second column shows the true income-emissions surfaces; the third column shows the estimated income-emissions surfaces (estimated functions plus estimated time fixed effects); and the fourth column shows function bias obtained as the difference between the surfaces in the second and third column. The results in Figure 12 can be compared to Figure 8, which contains estimation results from the same simulated data, but where we estimate the income-emissions relationship using the correct dynamic model.

Comparing Figures 8 and 12, it is clear the static model is not able to capture the time-varying income-emissions relationships. The implied, estimated surfaces from the static model (Figure 12) have a much higher bias than those of the dynamic model (Figure 8). This shows that if the income-emissions relationship is sufficiently varying in time, the static model with time fixed effects is not able to capture this. Neither would be popular alternatives such as the quadratic EKC model (Holtz-Eakin and Selden, 1995), the cubic EKC model (Grossman and Krueger, 1991), or the spline-based EKC model (Schmalensee et al., 1998), which also rely on time fixed effects. In case the income-emissions relationship is sufficiently varying in time, it is important to use a more flexible model such as the dynamic model proposed above.

5.4 Spuriousness

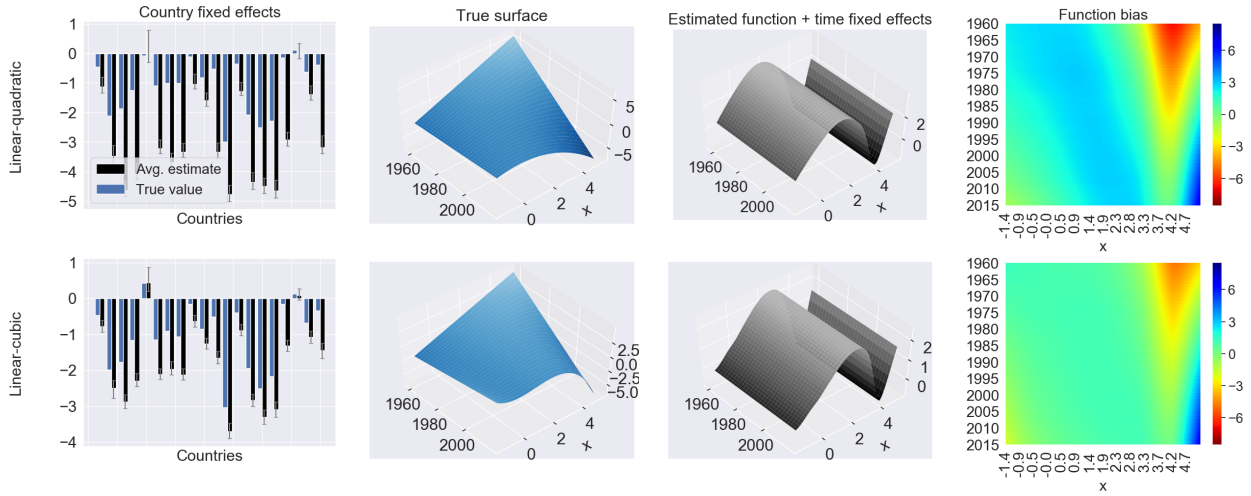Figure 13: Monte Carlo spuriousness results for the global formulation of the static model

Note: In the country fixed effects plot, gray error bars indicate 95% confidence intervals.

In this section, we investigate whether our estimator of the parameters in the static and the dynamic model suffers from spuriousness in a setting with independent but stochastically trending variables. Specifically, we simulate independent random walks to be used as the dependent variable in the models:

$$y_{it} = y_{it-1} + \nu_{it}, \quad i = 1, \ldots, N_t, \ t = 1, \ldots, T, \tag{5.1}$$

where $y_{i0} \overset{iid}{\sim} \mathcal{N}(0,1)$ and $\nu_{it} \overset{iid}{\sim} \mathcal{N}(0, \sigma_\nu^2)$. We then estimate a global $(R = 1)$ and a regional $(R = 5)$ formulation of both the static and the dynamic model using the methodology proposed above. That is, we estimate models where the dependent variable y_{it} is simulated from (5.1), and the input variable x_{it} is the natural logarithm of per capita GDP from the CO_2 emissions panel discussed in Section 3. Clearly, y_{it} and x_{it} are independent and characterized by country-specific stochastic trends. For ease of presentation, we again focus on results for the global model formulation. We report optimal network architectures based on an initial Monte Carlo sample in Table A.7 of the appendix.

Figure 14: Monte Carlo spuriousness results for the global formulation of the dynamic model
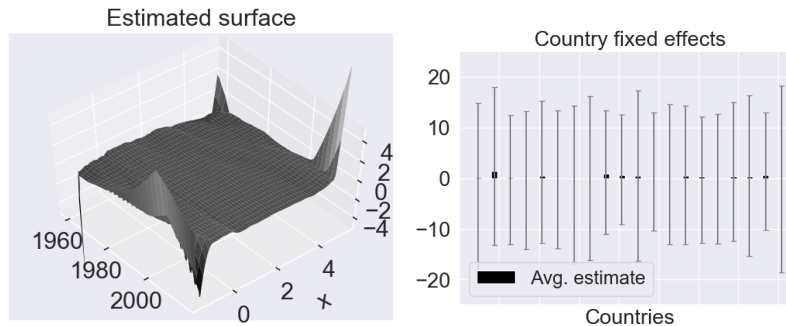
Note: In the country fixed effects plot, gray error bars indicate 95% confidence intervals.

Figures 13 and 14 show estimation results for the global formulation of the static and the dynamic model, respectively. For both models, the average estimate of the income-emissions relationship is around zero throughout the input region used for estimation, as it should be since the dependent variable and the input variable are completely unrelated. Similarly, average estimates of all fixed effects are around zero, as they should be. Results for the regional model formulation are similar. They are reported in Figures A.3 and A.4 of the appendix for the static and the dynamic model, respectively. The results of this section suggest our estimator does not suffer from spuriousness and can be reliably used in practice. The results are perhaps not surprising, as it is well-known that spuriousness is less problematic in panel data whenever the time dimension grows faster than the cross-sectional dimension (Levin et al., 2002).

6   Empirical Analysis

This section presents the results of our empirical analysis on the data introduced in Section 3. In Section 6.1, we use traditional territorial emissions, while Section 6.2 employs consumption-based emissions as the dependent variable.

6.1   Territorial emissions

To investigate the income-emissions relationship, we estimate a global and a regional formulation of both the static and the dynamic neural network model, as well as a national formulation of the static model. Optimal network architectures for each model, obtained by minimizing the BIC in (4.11), are presented in Table 4. As previously discussed, the learning task faced by the regional model formulation is more complex than that faced by its global counterpart, and similarly, the learning task faced by the national model formulation can be considered more complex than that faced by its regional counterpart. Likewise, the learning task faced by the dynamic model is in many ways more complex than that faced by its static counterpart. From Table 4, we see a clear tendency that increased complexity of the learning task prompts increased complexity of the neural network architecture. In particular, depth in the network architecture appears to be important for the dynamic model specification. The network architecture used for the regional formulation of the

Table 4: Optimal neural network architectures for territorial emissions

	Static model			Dynamic model	
	Global	Regional	National	Global	Regional
Network architecture	$(2,2)$	(8)	$(2,2,2)$	$(4,4,2)$	$(8,8,8)$
# parameters (excl. fixed effects)	12	56	388	44	208

Note: "(a,b,c)" indicates a neural network architecture with three hidden layers containing a units in the first layer, b in the second, and c in the third.

Table 5: Fractions of variance explained for territorial emissions

	Global	OECD	REF	Asia	MAF	LAM
Country effects only	0.92	0.79	0.86	0.81	0.91	0.85
Time effects only						
Global	0.04	-6.63	-0.52	-0.03	-0.23	0.05
Regional	0.38	0.04	0.07	0.11	0.05	0.10
Full static model						
Global	0.96	0.87	0.86	0.95	0.95	0.93
Regional	0.97	0.90	0.96	0.96	0.96	0.94
National	0.98	0.96	0.92	0.98	0.97	0.96
Full dynamic model						
Global	0.97	0.89	0.86	0.96	0.96	0.93
Regional	0.98	0.94	0.97	0.97	0.97	0.95

Note: Numbers in the table are R^2 statistics calculated using observations belonging to a given region. Global R^2 statistics are calculated using all observations. Row one is obtained by a regression using only country dummies. Row two is obtained by a regression using only time dummies. Row three is obtained by a regression in which time dummies have been interacted with a regional indicator. The last rows follow from estimation of the full models including income effects, country fixed effects, and potentially time fixed effects. We have highlighted the globally preferred formulation of each model.

dynamic model is the most complex. The national formulation of the static model may be regarded as the least parsimonious of the models considered, as it contains the largest number of parameters excluding fixed effects. However, simply comparing the number of parameters excluding fixed effects across the static and the dynamic model specification and across the three model formulations (global, regional and national) is misleading. For instance, the static model specification requires estimation of time fixed effects, while the dynamic model specification does not. Also, the national formulation of the static model contains many country-specific slope coefficients in the output layer, but does not require estimation of any fixed effects. Taking into account fixed effects, the regional formulation of the static model contains the most parameters, a total of 500. Likewise, when taking into account fixed effects, the global and the regional formulations of the static model contain more free parameters than their dynamic counterparts.

Table 5 provides information on how much sample variance is explained by the different model components, see the caption of the table for details. The first column of Table 5 shows how much of the global sample variance is explained, remaining columns show how much of the region-specific sample variance is explained by the respective models. We find that country fixed effects alone explain as much as 92% of the global sample variance, and that global time fixed effects explain only 4% of the global sample variance. So even though our data set spans six decades, cross-country

differences are more important than within-country differences. Similar conclusions were reached in Schmalensee et al. (1998). Note also that the fraction of variance explained varies across regions. In particular, we note a negative R^2 in the second row of Table 5 for the regions OECD, REF, Asia, and MAF, suggesting global time fixed effects fit these regions worse than a region-specific sample average. From the third row of Table 5, we see that allowing for region-specific time fixed effects helps increase the fraction of variance explained. This is true for all regions as well as globally. This is not surprising given that we found cross-country differences to be important. The last five rows of Table 5 show we can improve the fraction of variance explained by country fixed effects by between 6 and 17 percentage points, depending on the region, by specifying one of our proposed models. For both the static and the dynamic model specification, we note that a regional model formulation improves the fraction of variance explained by a global model formulation. This is true globally as well as for each region separately. Similarly, for the static model specification, we note that the national model formulation improves upon the fraction of variance explained by the regional model formulation, globally, as well as for each region separately except REF. We also note that for a given formulation (global or regional), the dynamic model never performs worse than the static model. Comparing the best performing formulations of the static and the dynamic model specifications, the national formulation of the static model and the regional formulation of the dynamic model appear to perform on par. Indeed, they explain the same amount of global sample variance to two decimals, and they explain the same amount of region-specific sample variance for the region MAF. The national formulation of the static model explains most region-specific sample variance for three out of the remaining four regions, and the regional formulation of the dynamic model explains most sample variance for the fourth region.

We show estimation results for the global and the regional formulation of the static neural network model in Figure 15. From the discussion in Section 5, recall that we are most confident about the model output within the interior of the input region used for estimation not too close to the boundaries. In Figure 15, we therefore shade the area between the .05 and the .95 income quantiles when we plot the estimated income-emissions relationship to indicate the area of highest confidence. The first row of plots in Figure 15 shows estimation results from the global model formulation. For reference, we also include estimation results from popular models within the EKC literature: a quadratic EKC model (Holtz-Eakin and Selden, 1995), a cubic EKC model (Grossman and Krueger, 1991), and a spline-based EKC model (Schmalensee et al., 1998). The benchmark models include country and time fixed effects like the static neural network model but employ different specifications of the emissions function. Remaining rows in Figure 15 show estimation results for the regional formulation of the static neural network model. For reference, we include results from region-wise estimation of the benchmark models. Note that region-wise estimation of the benchmark models ignores cross-region dependencies. By contrast, the regional formulation of the neural network models proposed in this paper uses information across regions to learn common input transformations.

Figure 15: Estimation results for the static model and territorial emissions
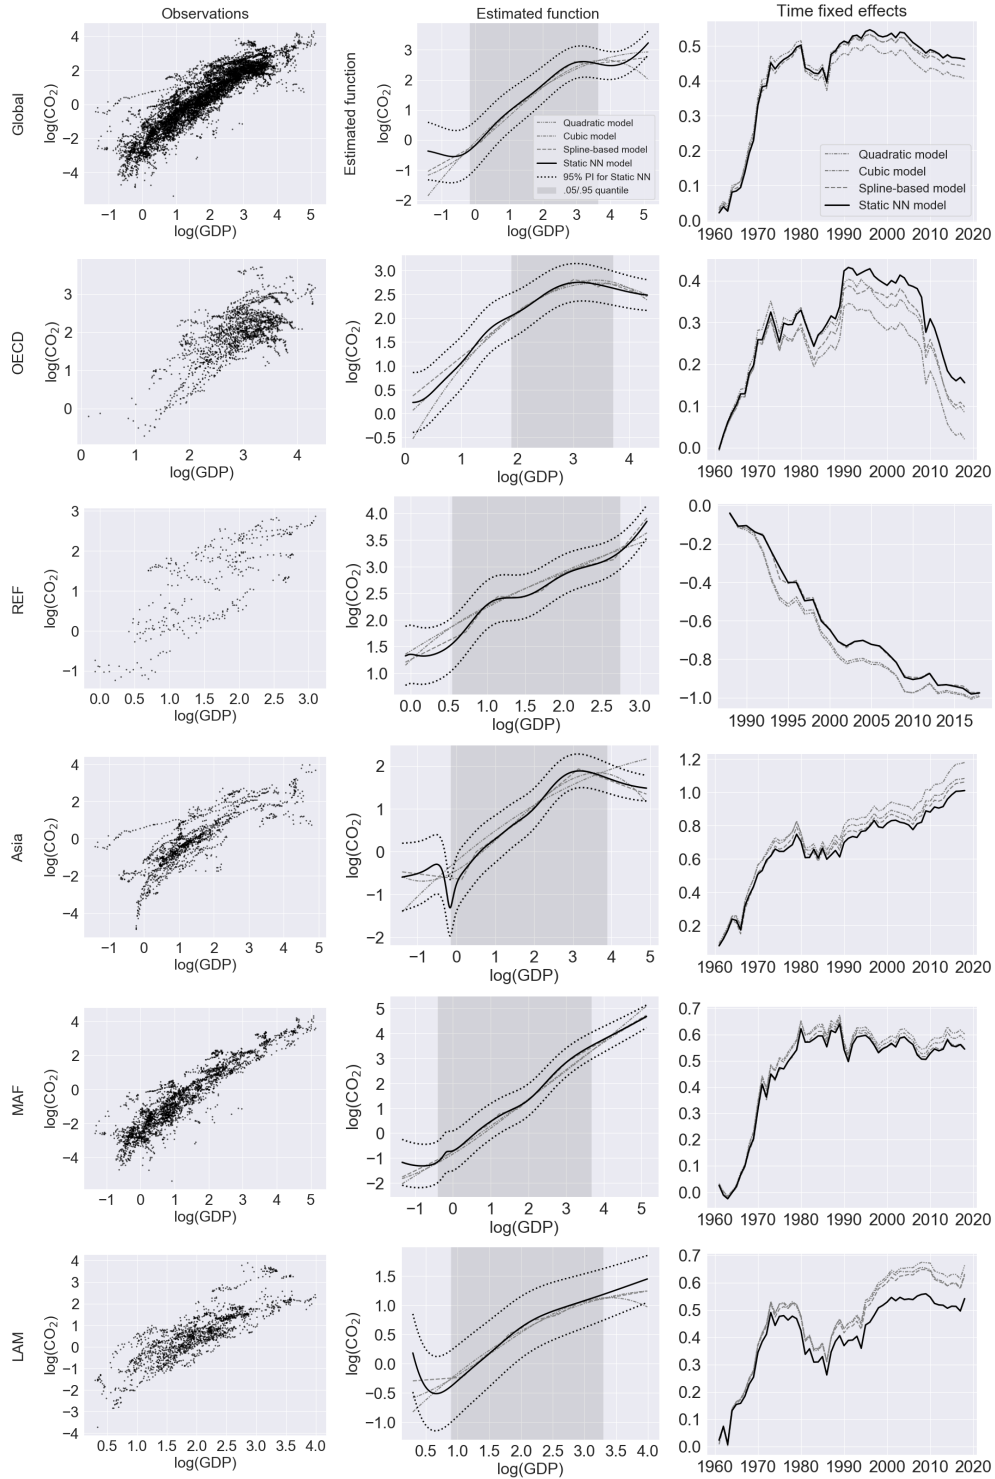
Note: Estimation results in the first row of plots are based on the global formulation of the static neural network model. Remaining rows are based on the regional formulation of the static neural network model. Benchmark models from the EKC literature are included in gray for reference.

Figure 16: Estimated income-emissions relationship for the world's nine largest CO_2 emitters based on territorial emissions
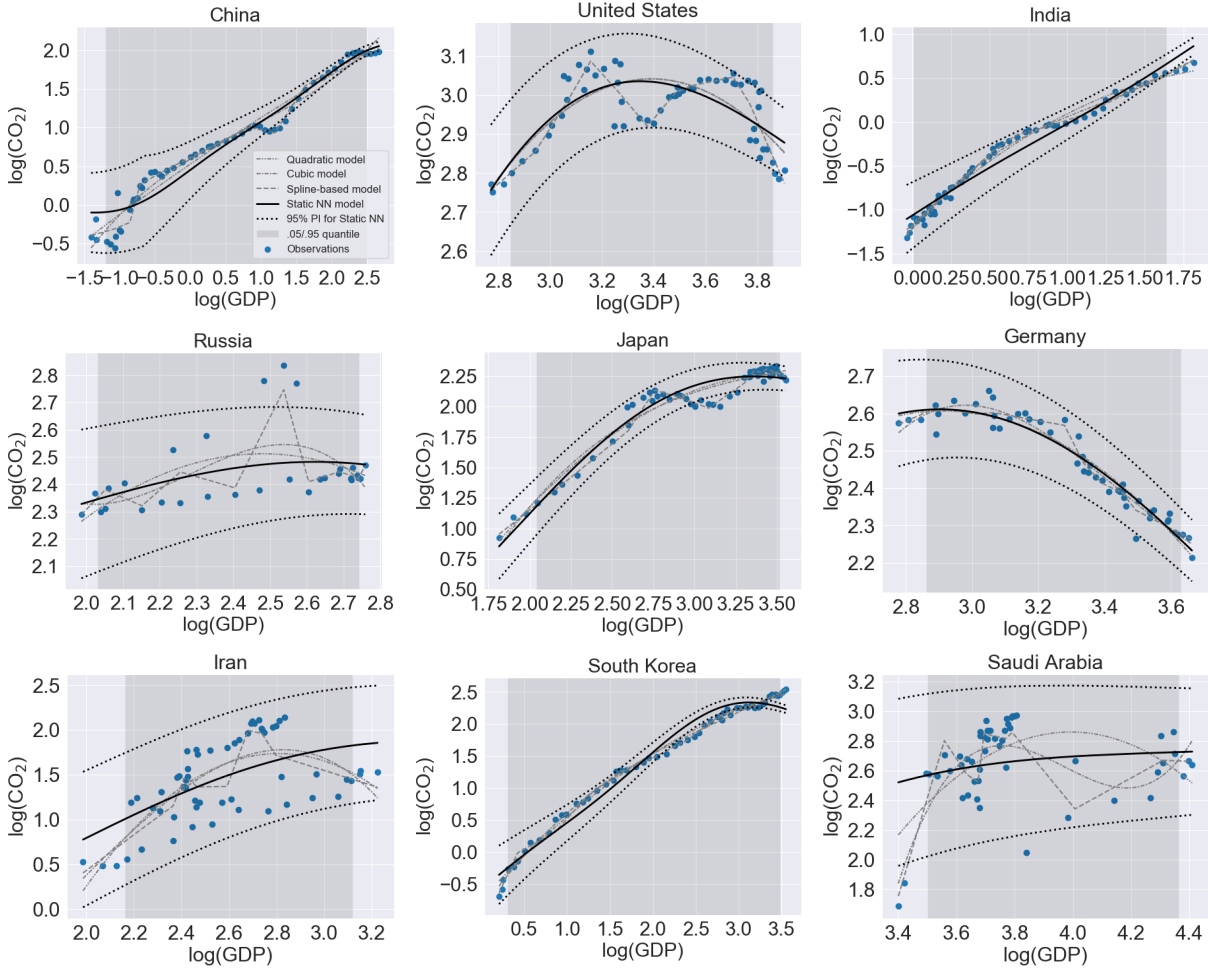
Note: Estimated income-emissions relationships are based on the national formulation of the static neural network model. Benchmark models from the EKC literature are included in gray for reference.

Inspecting Figure 15, we initially see evidence of a global EKC relationship for territorial CO_2 emissions. Indeed, for the global model formulation in the first row of the figure, we see evidence of an EKC-shape in both the estimated income-emissions relationship and in the estimated time fixed effects. From the regional model formulation, we also note evidence of an EKC-shape for OECD and Asia. For OECD, the static model shows a clear EKC-shape in both the estimated income-emissions relationship as well as in the estimated time fixed effects. For Asia, the static model again shows an EKC-shape in the estimated income-emissions relationship. For Asia, however, the EKC-shape is less clear in the estimated time fixed effects. It is interesting that we observe an EKC-shape for both OECD and Asia. Recall from the discussion in Section 3 that OECD is the richest of the five regions considered, while Asia is one of the poorest. Therefore, one could have expected that only countries within OECD would have become rich enough to reach a turning point. This does not seem to be the case. For the remaining regions REF, MAF, and LAM, the

Table 6: Country-specific shapes of the income-emissions relationship within each region for territorial emissions

	OECD	REF	Asia	MAF	LAM	Total
EKC	25	4	5	4	3	41
Increasing	10	6	29	53	28	126
Decreasing	6	3	1	5	3	18
Other	0	0	0	1	0	1

Note: The first row of the table reports the number of estimated EKC relationships within each region, defined as an inverse U-shaped relationship with an in-sample turning point, based on the national formulation of the static neural network model. The second and third rows of the table report the number of strictly increasing and decreasing estimated income-emissions relationship within each region, respectively. The fourth row of the table ("Other") represents Mozambique, which has an estimated U-shaped income-emissions relationship.

static model shows a monotonically increasing income-emissions relationship. The estimated time fixed effects display some concavity for these regions, and even a downward trend for the region REF. However, there is no evidence of an EKC-shape for these regions. Note the sample period does not start until 1990 for the region REF, and so the downward trending time fixed effects observed for this region in fact correspond to the period for which we also observe a downward trend in the time fixed effects for the regions OECD and MAF as well as for the global time fixed effects. Generally, results from the benchmark models appear in line with our static neural network model. But, globally and for the regions OECD and Asia, our static neural network model identifies a more pronounced EKC-shape and an earlier turning point.

We use the national formulation of the static model to estimate country-specific income-emissions relationships. Figure 16 shows estimates of the country-specific income-emissions relationships for the world's nine largest CO_2 emitting countries. For reference, we also include results from country-wise estimation of the benchmark models. By using cross-country information to learn common input transformations, it seems our static neural network model is generally able to learn more stable income-emissions relationships that are easier to interpret than the benchmark models. In Table 6, we report the number of estimated, country-specific EKCs within each region using our static neural network model, defined as an inverse U-shaped relationship with an in-sample turning point, as well as the number of strictly increasing, strictly decreasing, and other shapes of the income-emissions relationship. From Table 6, we see the most common shape across regions is the strictly increasing and the second most common shape is the EKC. We observe the most EKCs among OECD, which also has the largest ratio of EKCs to non-EKC-shapes. For Asia, the ratio of EKCs to non-EKC-shapes is considerably lower than for OECD. Yet, we observe a region-specific EKC for both of these regions when estimating a regional model formulation. Thus, it seems the regional EKC observed for Asia might be driven only by few wealthy countries such as Japan and South Korea, cf. Figure 16. Interestingly, the ratio of EKCs to non-EKC-shapes is

Figure 17: Estimation results for the dynamic model and territorial emissions
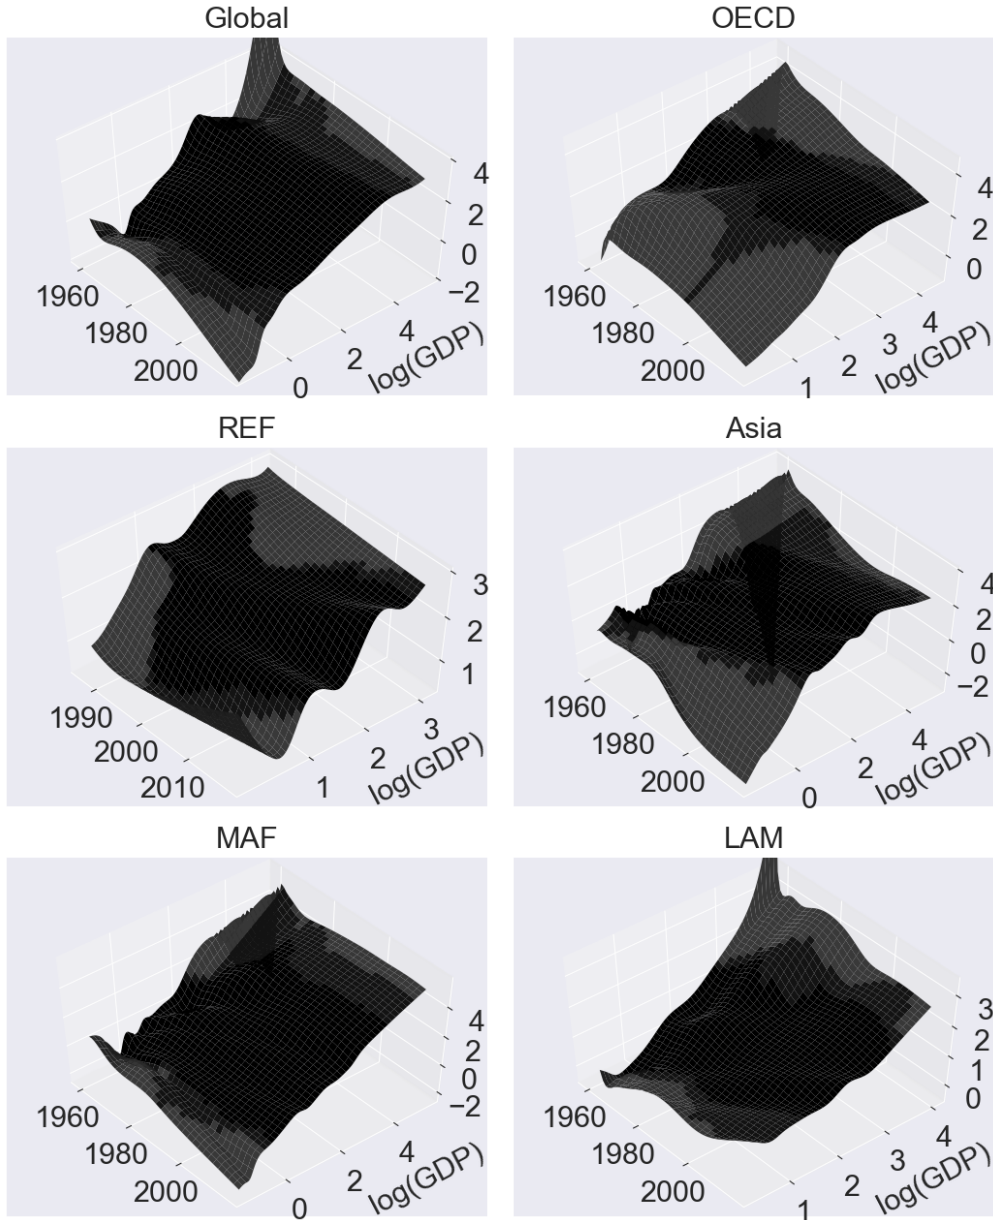
Note: Black coloring indicates area between the .05 and .95 income quantiles. Dark gray coloring indicates area with data support. Light gray coloring indicates area with no data support. The estimated surface in the top left panel is based on the global formulation of the dynamic neural network model. Remaining surfaces are based on the regional formulation of the dynamic neural network model.

also quite large for REF. However, this does not result in a region-specific EKC for REF when estimating a regional model formulation. This may be because the country-specific EKCs among REF are much less distinct than those among the OECD and Asia. The EKC for Russia in Figure 16 is representative of the EKCs among REF.

The dynamic neural network model allows us to take the analysis one step further. It allows

Figure 18: Time-varying environmental Kuznets curve for OECD based on territorial emissions
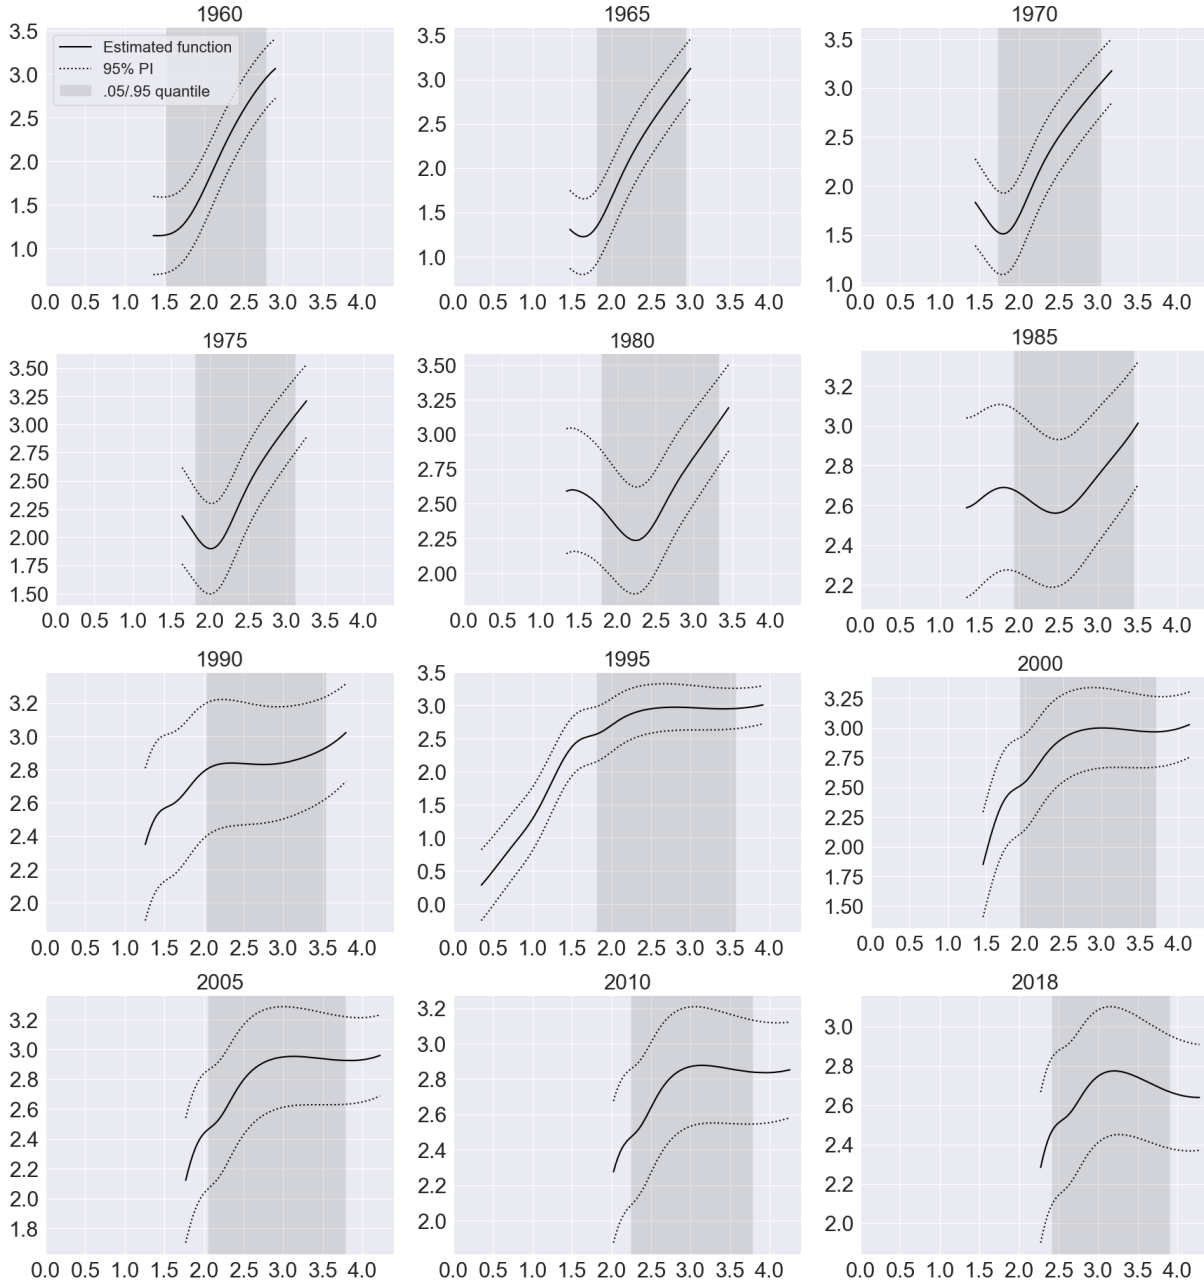
Note: Estimated income emissions relationships are based on the regional formulation of the dynamic neural network model.

us to investigate how the income-emissions relationships potentially change shape over time. We show estimation results for the dynamic neural network model in Figure 17. In the top left panel of the figure, we show the estimated surface from the global model formulation. Remaining panels show estimated surfaces from the regional model formulation. We indicate the area between the .05 and the .95 income quantiles by coloring the estimated surface black. Note that this area is

now varying with time. We also indicate the area outside of this time-varying quantile range but within the minimum and maximum income value observed for a given time period by coloring the estimated surface dark gray. The estimated surface is colored light gray over areas of the input space where the model has no data support.

Similar to the static model, the global formulation of the dynamic model shows evidence of a global EKC in the top left panel of Figure 17. Note it is straightforward to interpret the output of the dynamic model. There are no time fixed effects to take into account. We also note from Figure 17 that the global EKC-shape appears rather stable throughout the sample period. By contrast, it appears from the estimated surfaces using the regional formulation of the dynamic model that the EKC-shapes observed for OECD and Asia do not appear in the data until late in the sample period. Like the static model, the dynamic model also suggests a monotonically increasing income-emissions relationship for the regions REF, MAF, and LAM. The surface plot for the region REF is somewhat noisy, but, overall, there seems to be a monotonically increasing income-emissions relationship. The monotonically increasing relationships observed for the regions MAF and LAM appear rather stable over time.

In Figure 18, we illustrate how the income-emissions relationship for OECD changes shape over time by "slicing" the estimated surface for OECD every five years. We note from Figure 18 that the EKC-shape does not appear in the data for OECD until the early to mid 1990s. The turning point for OECD occurs around a value of log per capita GDP equal to 3, which is similar to what we find globally and for Asia in Figure 15, and it is supported by the benchmark models. However, it is interesting to note from Figure 18 that OECD reaches levels of log per capita GDP above 3 already in the 1960s, and log per capita GDP values above 3 is within the area between the .05 and .95 quantiles since the 1970s. Therefore, it seems a structural shift might occur around the early to mid 1990s that allows countries with high incomes to emit less CO_2. This shift could be the result of increased environmental concerns, climate change mitigation policies, exports, or other factors unrelated to income. For instance, the Montreal Protocol of the United Nations was adopted in 1987 with the goal of protecting the stratospheric ozone layer by phasing out production and consumption of ozone-depleting substances, and the United Nations Framework Convention on Climate Change (UNFCCC) was established in 1994. The shift could also reflect the pollution haven hypothesis that we discuss below.

6.2 Consumption-based emissions

We now investigate to what extend the income-emissions relationships, documented in the previous section, are driven, or affected, by international trade patterns. We do so by reestimating the global and the regional formulation of the static and the dynamic neural network model using consumption-based emissions for the dependent variable as opposed to territorial emissions; see Section 3 for the definition of consumption-based emissions. Since consumption-based emissions are available only from year 1990 onward, leaving the time series available for each country rather short, we do not estimate a national formulation of the static model. To facilitate comparison of the results obtained from using territorial and consumption-based emissions, we also reestimate all models for territorial emissions restricting the sample to coincide with the sample for consumption-based

37

Table 7: Optimal neural network architectures for consumption-based emissions

	Static model		Dynamic model	
	Global	Regional	Global	Regional
$CO_2{}^C$				
Neural network architecture	$(2, 2)$	(4)	$(4, 2)$	$(8, 4)$
# parameters (excl. fixed effects)	12	28	24	80
$CO_2{}^\star$				
Neural network architecture	$(2, 2)$	(4)	$(4, 2, 2)$	$(8, 8, 8)$
# parameters (excl. fixed effects)	12	28	30	208

Note: "$CO_2{}^C$" are consumption-based CO_2 emissions; "$CO_2{}^\star$" are territorial CO_2 emissions based on a restricted sample that coincides with the one for consumption-based emissions; "(a,b,c)" indicates a neural network architecture with three hidden layers containing a units in the first layer, b in the second, and c in the third.

emissions, i.e. beginning in 1990. This shows whether observed changes in the income-emissions relationships can plausibly be attributed to the switch from territorial emission to consumption-based emissions, or whether changes are simply an effect of the restricted data set.

In the first row of Table 7, we report optimal neural network architectures for consumption-based emissions, obtained by minimizing the BIC (4.11). In the third row of Table 7, we report optimal neural network architectures for territorial emissions using the restricted sample. Overall, Table 7 is broadly comparable to Table 4 of the previous section. Namely, a more complex network architecture is required for the dynamic model than for the static model, and a more complex network architecture is required for a regional model formulation than for a global model formulation. Comparing the first and the third row of Table 7, we also note that, for the dynamic model, a higher degree of model complexity is required for territorial emissions on the restricted sample than for consumption-based emissions.

Figure 19 shows estimation results for the static neural network model for consumption-based emissions as well as territorial emissions based on the restricted sample. The first row of plots in the figure shows estimation results for the global formulation of the static model. Remaining rows show estimation results for the regional formulation of the static model. In plots of the estimated income-emissions relationship (second column) and time fixed effects (third column), a solid line indicates results for consumption-based emissions, and a dashed line indicates results for territorial emissions based on the restricted sample. Estimation results for the quadratic, the cubic, and the spline-based benchmark models using consumption-based emissions for the dependent variable can be found in Figure A.5 of the appendix. They are generally in line with the results from the static neural network model in Figure 19.

Inspecting the results from the global formulation of the static model in the first row of Figure 19, we note a downward trend in the estimated time fixed effects for territorial emissions. However, recall the sample does not start until year 1990. This implies that the downward trend in the esti-

Figure 19: Estimation results for the static model and consumption-based emissions
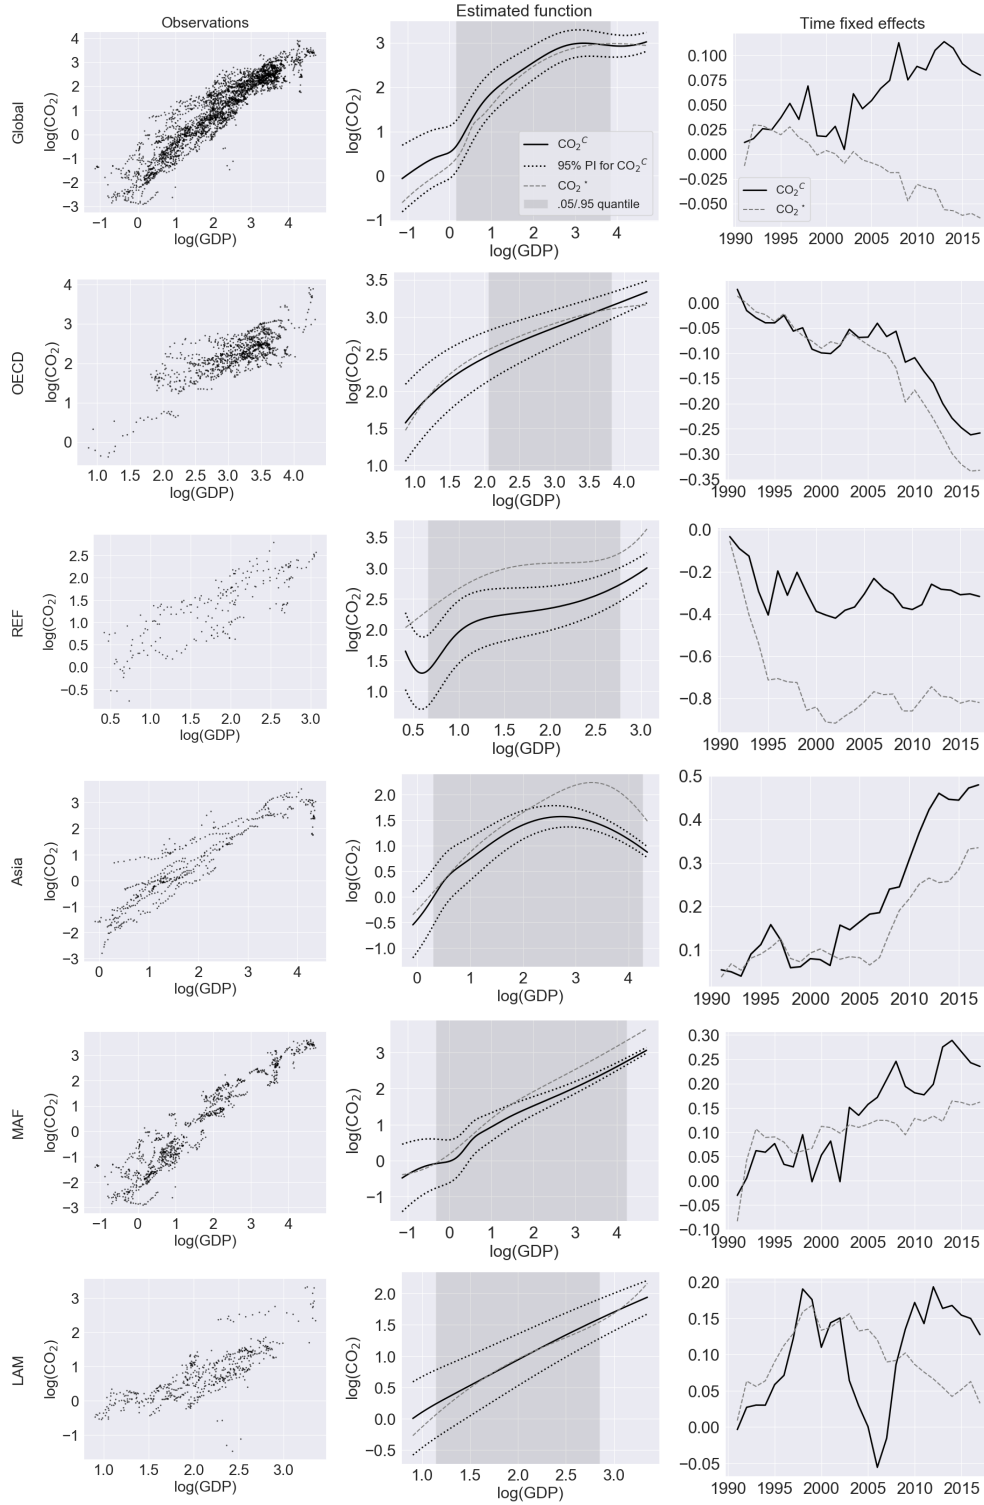
Note: "$CO_2{}^C$" are consumption-based CO_2 emissions and "$CO_2{}^\star$" are territorial CO_2 emissions based on a restricted sample that coincides with the one for consumption-based emissions. Estimation results in the first row of plots are based on the global formulation of the static neural network model. Remaining rows are based on the regional formulation of the static neural network model.

mated time fixed effects for territorial emissions in the top right panel of Figure 19 corresponds to the downward trend in estimated time fixed effects for that same period in the top right panel of Figure 15. Conversely, estimated time fixed effects for consumption-based emissions are trending upwards in the top right panel of Figure 19. But, we also note that the magnitude of the estimated time fixed effects over this period is small, suggesting the effect of time on the global income-emission relationship is rather minimal over this sample period. In the top middle panel of Figure 19, we note the estimated income-emissions relationship from the global model formulation exhibits concavity for both types of emissions in line with the EKC hypothesis. However, the global EKC-shape is not quite as pronounced as in Figure 15. We conclude that when controlling for international trade flows, we still observe concavity in the global income-emissions relationship, i.e. the EKC relationship appears intact.

Inspecting the results from the regional formulation of the static model in the remaining rows of Figure 19, we immediately note the previously observed EKC-shape for OECD when using territorial emissions disappears when using consumption-based emissions. The income-emissions relationship appears to be strictly increasing for OECD when using consumption-based emissions and thus not reaching a turning point or even flattening out. Although the EKC-shape for OECD is not quite as pronounced for territorial emissions based on the restricted sample as it was for the full sample in Figure 15, there still is concavity for territorial emissions based on the restricted sample. Recall from the discussion in Section 3 that OECD is the only region that is a net-exporter of CO_2 emissions to other countries throughout the restricted sample period from year 1990 onward. Together with the fact that the EKC-shape disappears for OECD when using consumption-based emissions, this suggests that the apparent EKC relationship we found when using territorial emissions can be plausibly explained by the pollution haven hypothesis. Recall the pollution haven hypothesis suggests that rich countries, in this case belonging to the OECD region, export their emissions to other countries, often developing countries subject to weak environmental regulation, through international trade. For Asia, we note, by contrast, that the EKC-shape becomes even crisper and with an earlier turning point when using consumption-based emissions than when using territorial emissions. We interpret this as the mirror image of the result for OECD and as further support for the pollution haven hypothesis; recall also from the discussion in Section 3 that Asia is a net importer of CO_2 emissions throughout the restricted sample period. For most of the remaining regions, the shape of the estimated income-emissions relationship does not change substantially when moving from territorial emissions to consumption-based emissions. For REF, there is a substantial level shift in the income-emissions relationship estimated by the static model when moving from territorial emissions to consumption-based emissions. However, we observe a similar level shift in the estimated time fixed effects, and so this level shift is most likely the result of an identification issue stemming from the small number of observations available for this region.

We show estimation results from the global and the regional formulation of the dynamic neural network model when using consumption-based emission for the dependent variable in Figure 20. Estimation results from the dynamic model when using territorial emissions based on the restricted sample for the dependent variable can be found in Figure A.6 of the appendix; they look similar to the results from the dynamic model when using territorial emissions based on the unrestricted

Figure 20: Estimation results for the dynamic model and consumption-based emissions
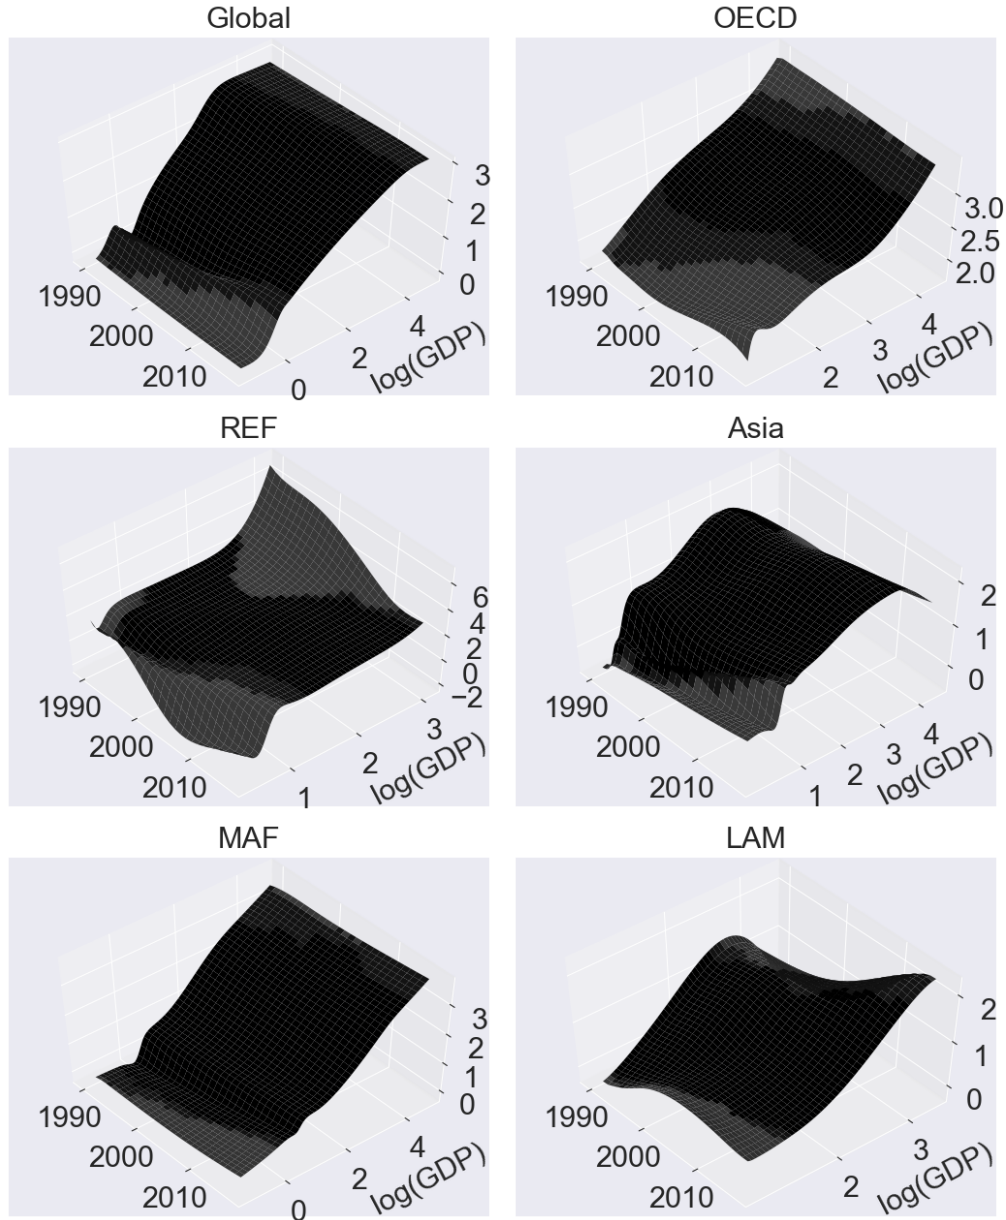
Black coloring indicates area between the .05 and .95 income quantiles. Dark gray coloring indicates area with data support. Light gray coloring indicates area with no data support. The estimated surface in the top left panel is based on the global formulation of the dynamic neural network model. Remaining surfaces are based on the regional formulation of the dynamic neural network model.

sample for the dependent variable in Figure 17. Inspecting the estimated surface in the top left panel of Figure 20 obtained from the global model formulation, it shows concavity in the income dimension, and it seems stable in the time dimension. As such, the conclusion from the global formulation of the static model of a global EKC also for consumption-based emissions seems stable over the restricted sample. Inspecting the estimated surfaces from the regional model formulation

41

in the remaining panels of Figure 20, we note also the conclusions from the regional formulation of the static model appear stable over the restricted sample. In particular, the strictly increasing income-emissions relationship observed for OECD when using consumption-based emissions for the dependent variable appear stable over the restricted sample, and the EKC-shape observed for Asia when using consumption-based emissions for the dependent variable appear stable over the restricted sample.

7 Conclusion

We propose a novel neural network-based panel data methodology for analyzing the environmental Kuznets curve (EKC) for carbon dioxide (CO_2) emissions. We consider two distinct model specifications within this overall framework: a static model specification and a dynamic model specification. The static model consists of both country and time fixed effects in addition to a feedforward neural network component with income as the only input variable. This model is static in the sense that the shape of the income-emissions relationship is assumed to be fixed over time. The dynamic model uses a time variable as an additional input into the neural network component in place of time fixed effects. By using both income and a time variable as inputs into the neural network component, the dynamic model is able to learn how time and income interact, and how the income-emissions relationship potentially changes its entire shape over time. Both model specifications use cross-country dependencies to learn common input transformations by having some model parameters be shared across all countries, and simultaneously allow for cross-country heterogeneity in the shape of the income-emissions relationship by having other parameters be specific to regions of homogeneous countries. In a Monte Carlo study, we demonstrate that our proposed methodology is able to identify various functional forms of different complexity, and we also demonstrate its ability to account for country-specific stochastic trends. Furthermore, we demonstrate that the dynamic model is able to capture time-varying income-emissions relationships that cannot be captured using time fixed effects.

We investigate the relationship between per capita GDP and CO_2 emissions, using national-level panel data for the period 1960-2018. When using territorial emissions data that measures emissions based on production, we find that globally as well as for the regions OECD and Asia, there is evidence of an inverse U-shaped income-emissions relationship, often referred to as an EKC. The dynamic model suggests the global relationship is rather stable over time. However, for both OECD and Asia, it seems the EKC-shape does not appear until late in the sample period. Conversely, when using consumption-based emissions data that accounts for cross-country emissions transfers through international trade, the evidence of an EKC relationship for OECD disappears, while, for Asia, the EKC-relationship becomes clearer and with an earlier turning point. This suggests the apparent EKC-relationship observed for OECD when using territorial emissions is driven by emissions exports to other countries. It also suggests that, for Asia, there are EKC-effects in emissions due to local consumption that are not seen in territorial emissions due to imported emissions.

Acknowledgements

References

Agras, J. and D. Chapman (1999). A dynamic approach to the Environmental Kuznets Curve hypothesis. *Ecological Economics 28*(2), 267 – 277.

Anders, U. and O. Korn (1999). Model selection in neural networks. *Neural Networks 12*(2), 309 – 323.

Arrow, K., B. Bolin, R. Costanza, P. Dasgupta, C. Folke, C. Holling, B.-O. Jansson, S. Levin, K.-G. Mäler, C. Perrings, and D. Pimentel (1995). Economic growth, carrying capacity, and the environment. *Ecological Economics 15*(2), 91 – 95.

Aslanidis, N. (2009, September). Environmental Kuznets Curves for Carbon Emissions: A Critical Survey. Working Papers 2009.75, Fondazione Eni Enrico Mattei.

Auffhammer, M. and R. Steinhauser (2012). Forecasting The Path of U.S. CO2 Emissions Using State-Level Information. *The Review of Economics and Statistics 94*(1), 172–185.

Azomahou, T., F. Laisney, and P. N. Van (2006). Economic development and CO2 emissions: A nonparametric panel approach. *Journal of Public Economics 90*(6), 1347 – 1363.

Barros, F. G., A. F. Mendonça, and J. M. Nogueira (2002). Poverty and environmental degradation: The Kuznets environmental curve for the Brazilian case. *Universidade de Brasília* (267).

Bertinelli, L. and E. Strobl (2005). The Environmental Kuznets Curve semi-parametrically revisited. *Economics Letters 88*(3), 350 – 357.

Blundell, C., J. Cornebise, K. Kavukcuoglu, and D. Wierstra (2015). Weight uncertainty in neural networks.

Borghesi, S. (2006). Income inequality and the environmental Kuznets curve. *Environment, inequality and collective action 33*.

Cole, M. A. (2003). Development, trade, and the environment: how robust is the Environmental Kuznets Curve? *Environment and Development Economics 8*(4), 557–580.

Cole, M. A. (2004). Trade, the pollution haven hypothesis and the environmental Kuznets curve: examining the linkages. *Ecological Economics 48*(1), 71 – 81.

Costantini, V. and C. Martini (2006). A modified environmental Kuznets curve for sustainable development assessment using panel data.

Cybenko, G. (1989, Dec). Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems 2*(4), 303–314.

Dasgupta, S., K. Hamilton, K. D. Pandey, and D. Wheeler (2006). Environment During growth: Accounting for governance and vulnerability. *World Development 34*(9), 1597 – 1611.

de Bruyn, S., J. van den Bergh, and J. Opschoor (1998). Economic growth and emissions: reconsidering the empirical basis of environmental Kuznets curves. *Ecological Economics 25*(2), 161 – 175.

Dijkgraaf, E. and H. R. J. Vollebergh (2005, October). A Test for Parameter Homogeneity in CO2 Panel EKC Estimations. *Environmental and Resource Economics 32*(2), 229—239.

Dinda, S. (2004). Environmental Kuznets Curve Hypothesis: A Survey. *Ecological Economics 49*(4), 431 – 455.

Dutt, K. (2009, August). Governance, institutions and the environment-income relationship: a cross-country study. *Environment, Development and Sustainability 11*(4), 705—723.

Galeotti, M., A. Lanza, and F. Pauli (2006). Reassessing the environmental Kuznets curve for CO2 emissions: A robustness exercise. *Ecological Economics 57*(1), 152 – 163.

Global Carbon Project (2019). Supplemental data of Global Carbon Budget 2019 (Version 1.0) [Data set]. Global Carbon Project.

Glorot, X., A. Bordes, and Y. Bengio (2011, 11–13 Apr). Deep Sparse Rectifier Neural Networks. In G. Gordon, D. Dunson, and M. Dudík (Eds.), *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, Volume 15 of *Proceedings of Machine Learning Research*, Fort Lauderdale, FL, USA, pp. 315–323. PMLR.

Goodfellow, I., Y. Bengio, and A. Courville (2016). *Deep Learning*. MIT Press. http://www.deeplearningbook.org.

Grossman, G. M. and A. B. Krueger (1991, November). Environmental Impacts of a North American Free Trade Agreement. Working Paper 3914, National Bureau of Economic Research.

Gu, S., B. Kelly, and D. Xiu (2020, 02). Empirical Asset Pricing via Machine Learning. *The Review of Financial Studies 33*(5), 2223–2273.

Halkos, G. E. and E. G. Tsionas (2001). Environmental Kuznets curves: Bayesian evidence from switching regime models. *Energy Economics 23*(2), 191 – 210.

Harbaugh, W. T., A. Levinson, and D. M. Wilson (2002). Reexamining the Empirical Evidence for an Environmental Kuznets Curve. *The Review of Economics and Statistics 84*(3), 541–551.

He, K., X. Zhang, S. Ren, and J. Sun (2015, December). Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. In *The IEEE International Conference on Computer Vision (ICCV)*.

Hendry, D. (2018). First-in, first-out: Driving the UK's per capita carbon dioxide emissions below 1860 levels. *VoxEU*.

Hill, R. J. and E. Magnani (2002). An Exploration of the Conceptual and Empirical Basis of the Environmental Kuznets Curve. *Australian Economic Papers 41*(2), 239–254.

Holtz-Eakin, D. and T. M. Selden (1995). Stoking the fires? CO2 emissions and economic growth. *Journal of Public Economics 57*(1), 85 – 101.

Hornik, K., M. Stinchcombe, and H. White (1989). Multilayer feedforward networks are universal approximators. *Neural Networks 2*(5), 359–366.

IPCC (2014). AR5 Synthesis Report: Climate Change 2014.

Kearsley, A. and M. Riddel (2010). A further inquiry into the Pollution Haven Hypothesis and the Environmental Kuznets Curve. *Ecological Economics 69*(4), 905 – 919. Special Section: Coevolutionary Ecological Economics: Theory and Applications.

Kingma, D. P. and J. Ba (2014). Adam: A Method for Stochastic Optimization.

Kuznets, S. (1955). Economic Growth and Income Inequality. *The American Economic Review 45*(1), 1–28.

Lee, C.-C. and J.-D. Lee (2009). Income and CO2 emissions: Evidence from panel unit root and cointegration tests. *Energy Policy 37*(2), 413 – 423.

Leshno, M., V. Y. Lin, A. Pinkus, and S. Schocken (1993). Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural Networks 6*(6), 861–867.

Levin, A., C.-F. Lin, and C.-S. J. Chu (2002). Unit root tests in panel data: asymptotic and finite-sample properties. *Journal of Econometrics 108*(1), 1–24.

List, J. A. and C. A. Gallet (1999). The environmental Kuznets curve: does one size fit all? *Ecological Economics 31*(3), 409 – 423.

Magnani, E. (2001). The Environmental Kuznets Curve: development path or policy result? *Environmental Modelling & Software 16*(2), 157 – 165. Environmental Modelling and Socioeconomics.

Malec, P. and M. Schienle (2014). Nonparametric kernel density estimation near the boundary. *Computational Statistics & Data Analysis 72*, 57 – 76.

Masters, T. (1993). *Practical neural network recipes in C++*. (Academic Press).

Millimet, D. L., J. A. List, and T. Stengos (2003). The Environmental Kuznets Curve: Real Progress or Misspecified Models? *The Review of Economics and Statistics 85*(4), 1038–1047.

Mundlak, Y. (1978). On the Pooling of Time Series and Cross Section Data. *Econometrica 46*(1), 69–85.

Müller-Fürstenberger, G. and M. Wagner (2007). Exploring the environmental Kuznets hypothesis: Theoretical and econometric problems. *Ecological Economics 62*(3), 648 – 660.

Nix, D. and A. Weigend (1994a). Learning local error bars for nonlinear regression. In G. Tesauro, D. Touretzky, and T. Leen (Eds.), *Advances in Neural Information Processing Systems*, Volume 7. MIT Press.

Nix, D. A. and A. S. Weigend (1994b). Estimating the mean and variance of the target probability distribution. *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94) 1*, 55–60 vol.1.

NOAA (2020). The NOAA Annual Greenhouse Gas Index. Accessed June 2020. https://www.esrl.noaa.gov/gmd/aggi/aggi.html.

Panayotou, T. (1997). Demystifying the environmental Kuznets curve: turning a black box into a policy tool. *Environment and Development Economics 2*(4), 465–484.

Perman, R. and D. Stern (1999). The Environmental Kuznets Curve: Implications of Non-Stationarity. Working papers in ecological economics, Australian National University, Centre for Resource and Environmental Studies, Ecological Economics Program.

Perman, R. and D. I. Stern (2003). Evidence from panel unit root and cointegration tests that the Environmental Kuznets Curve does not exist. *Australian Journal of Agricultural and Resource Economics 47*(3), 325–347.

Peters, G. P. and E. G. Hertwich (2008a, March). CO2 Embodied in International Trade with Implications for Global Climate Policy. *Environmental Science & Technology 42*(5), 1401—1407.

Peters, G. P. and E. G. Hertwich (2008b, January). Post-Kyoto greenhouse gas inventories: production versus consumption. *Climatic Change 86*(1), 51—66.

Peters, G. P., J. C. Minx, C. L. Weber, and O. Edenhofer (2011). Growth in emission transfers via international trade from 1990 to 2008. *Proceedings of the National Academy of Sciences 108*(21), 8903–8908.

Ramachandran, P., B. Zoph, and Q. V. Le (2017). Searching for Activation Functions.

Riahi, K., D. P. van Vuuren, E. Kriegler, J. Edmonds, B. C. O'Neill, S. Fujimori, N. Bauer, K. Calvin, R. Dellink, O. Fricko, W. Lutz, A. Popp, J. C. Cuaresma, S. KC, M. Leimbach, L. Jiang, T. Kram, S. Rao, J. Emmerling, K. Ebi, T. Hasegawa, P. Havlik, F. Humpenöder, L. A. D. Silva, S. Smith, E. Stehfest, V. Bosetti, J. Eom, D. Gernaat, T. Masui, J. Rogelj, J. Strefler, L. Drouet, V. Krey, G. Luderer, M. Harmsen, K. Takahashi, L. Baumstark, J. C. Doelman, M. Kainuma, Z. Klimont, G. Marangoni, H. Lotze-Campen, M. Obersteiner, A. Tabeau, and M. Tavoni (2017). The shared socioeconomic pathways and their energy, land use, and greenhouse gas emissions implications: An overview. *Global Environmental Change 42*, 153 – 168.

Richmond, A. K. and R. K. Kaufmann (2006, February). Is there a turning point in the relationship between income and energy use and/or carbon emissions? *Ecological Economics 56*(2), 176–189.

Romero-Ávila, D. (2008). Questioning the empirical basis of the environmental Kuznets curve for CO2: New evidence from a panel stationarity test robust to multiple breaks and cross-dependence. *Ecological Economics 64*(3), 559 – 574.

Schmalensee, R., T. Stoker, and R. Judson (1998, 02). World Carbon Dioxide Emissions: 1950-2050. *The Review of Economics and Statistics 80*, 15–27.

Selden, T. M. and D. Song (1994). Environmental Quality and Development: Is There a Kuznets Curve for Air Pollution Emissions? *Journal of Environmental Economics and Management 27*(2), 147 – 162.

Shafik, N. (1994). Economic Development and Environmental Quality: An Econometric Analysis. *Oxford Economic Papers 46*, 757–773.

Stern, D. and M. Common (2001, 02). Is There An Environmental Kuznets Curve for Sulphur. *Journal of Environmental Economics and Management 41*, 162–178.

Stern, D. I. (2004). The Rise and Fall of the Environmental Kuznets Curve. *World Development 32*(8), 1419 – 1439.

Stern, D. I. (2010). Between estimates of the emissions-income elasticity. *Ecological Economics 69*(11), 2173 – 2182. Special Section - Payments for Ecosystem Services: From Local to Global.

Stern, D. I. (2017, Apr). The environmental Kuznets curve after 25 years. *Journal of Bioeconomics 19*(1), 7–28.

Stern, D. I., M. S. Common, and E. B. Barbier (1996). Economic growth and environmental degradation: The environmental Kuznets curve and sustainable development. *World Development 24*(7), 1151 – 1160.

Suri, V. and D. Chapman (1998). Economic growth, trade and energy: implications for the environmental Kuznets curve. *Ecological Economics 25*(2), 195 – 208.

Taskin, F. and O. Zaim (2000). Searching for a Kuznets curve in environmental efficiency using kernel estimation. *Economics Letters 68*(2), 217 – 223.

Torras, M. and J. K. Boyce (1998). Income, inequality, and pollution: a reassessment of the environmental Kuznets Curve. *Ecological Economics 25*(2), 147 – 160.

Tsurumi, T. and S. Managi (2015). *Environmental Kuznets curve: Economic growth and emission reduction.* cited By 3.

Vollebergh, H. R., B. Melenberg, and E. Dijkgraaf (2009). Identifying reduced-form relations with panel data: The case of pollution and income. *Journal of Environmental Economics and Management 58*(1), 27 – 42.

Wagner, M. (2008). The carbon Kuznets curve: A cloudy picture emitted by bad econometrics? *Resource and Energy Economics 30* (3), 388 – 408.

Wagner, M. (2015). The Environmental Kuznets Curve, Cointegration and Nonlinearity. *Journal of Applied Econometrics 30* (6), 948–967.

Zapranis, A. and E. Livanis (2005). Prediction intervals for neural network models. In *Proceedings of the 9th WSEAS International Conference on Computers*, ICCOMP'05, Stevens Point, Wisconsin, USA. World Scientific and Engineering Academy and Society (WSEAS).

Appendix

A.1 Figures

Figure A.1: Monte Carlo misspecification results for the regional formulation of the static model: linear-quadratic function
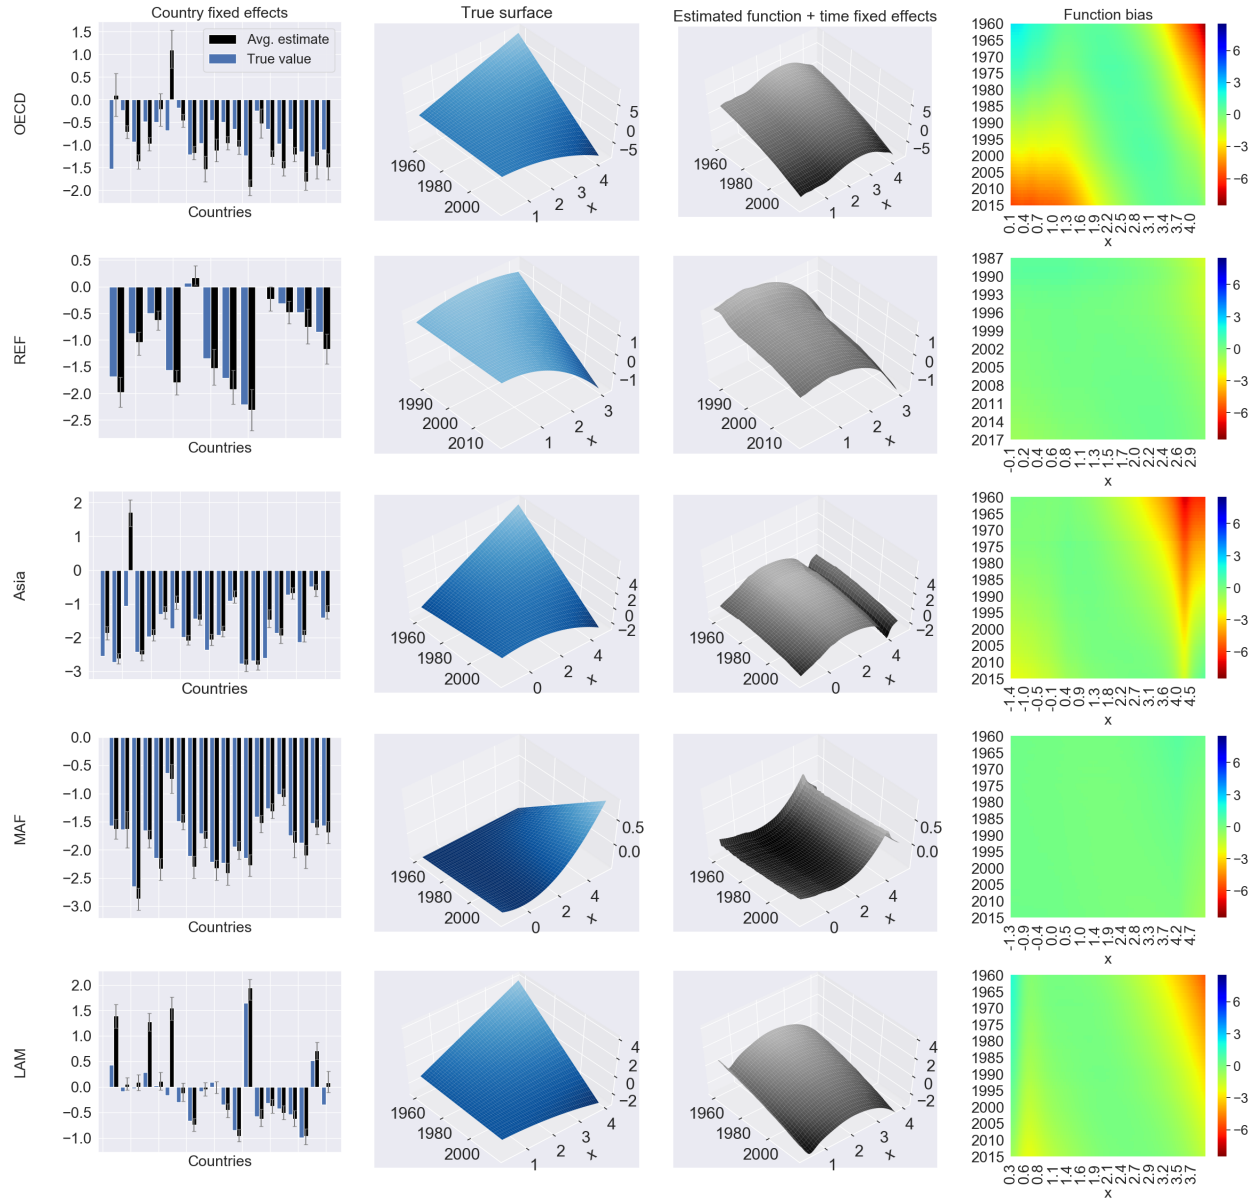
Note: In the country fixed effects plots, gray error bars indicate 95% confidence intervals.

Figure A.2: Monte Carlo misspecification results for the regional formulation of the static model: linear-cubic function
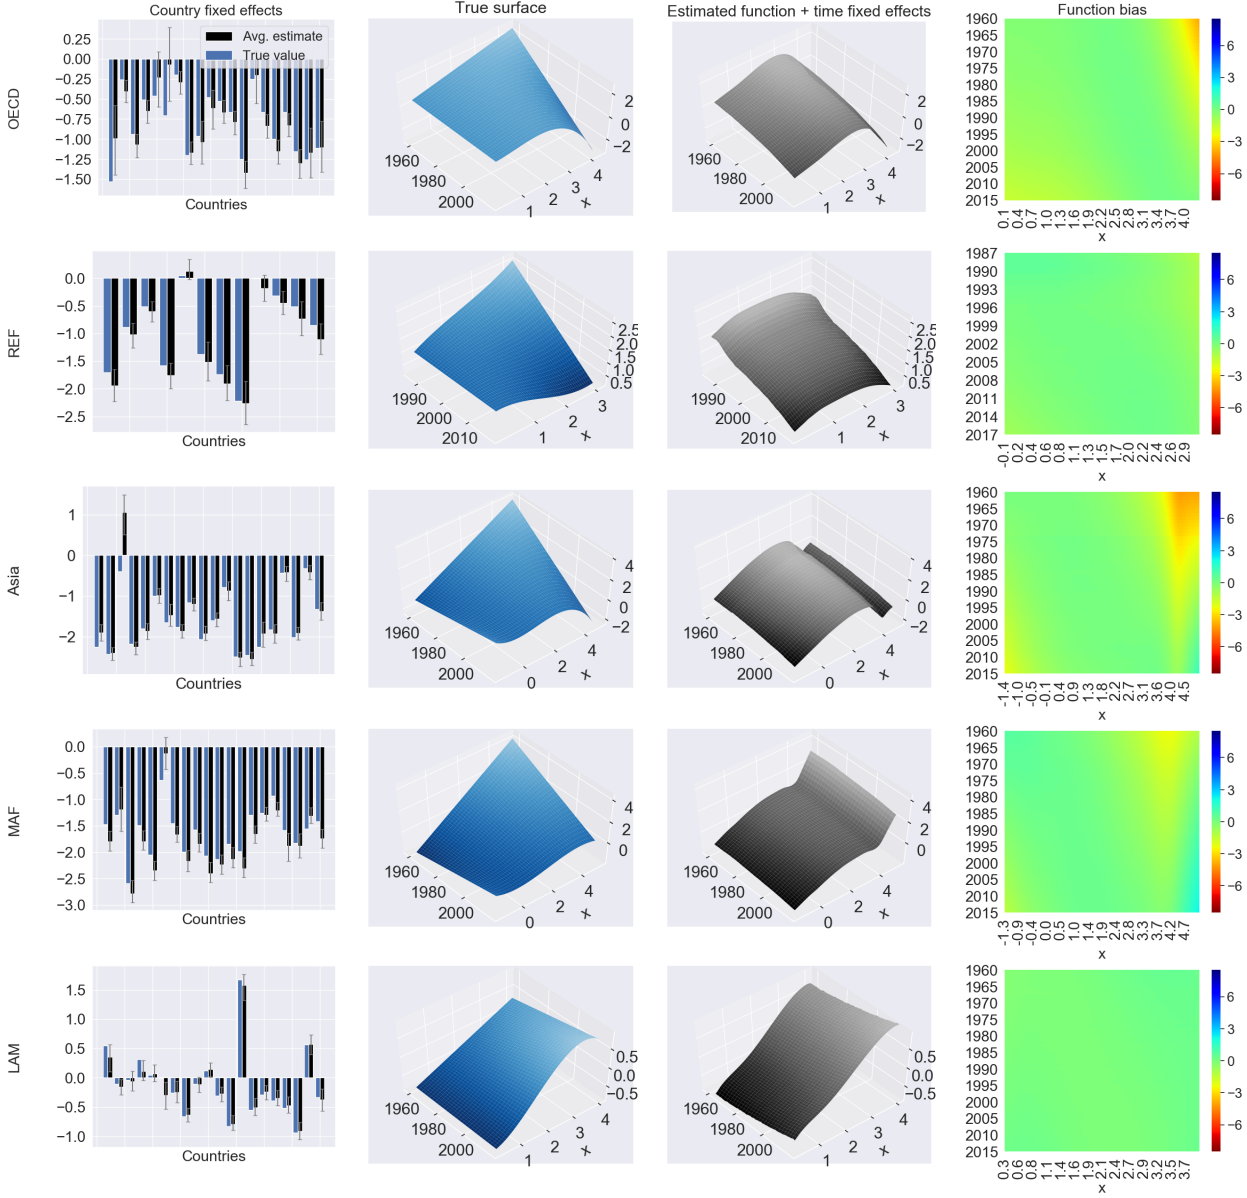
Note: In the country fixed effects plots, gray error bars indicate 95% confidence intervals.

Figure A.3: Monte Carlo spuriousness results for the regional formulation of the static model
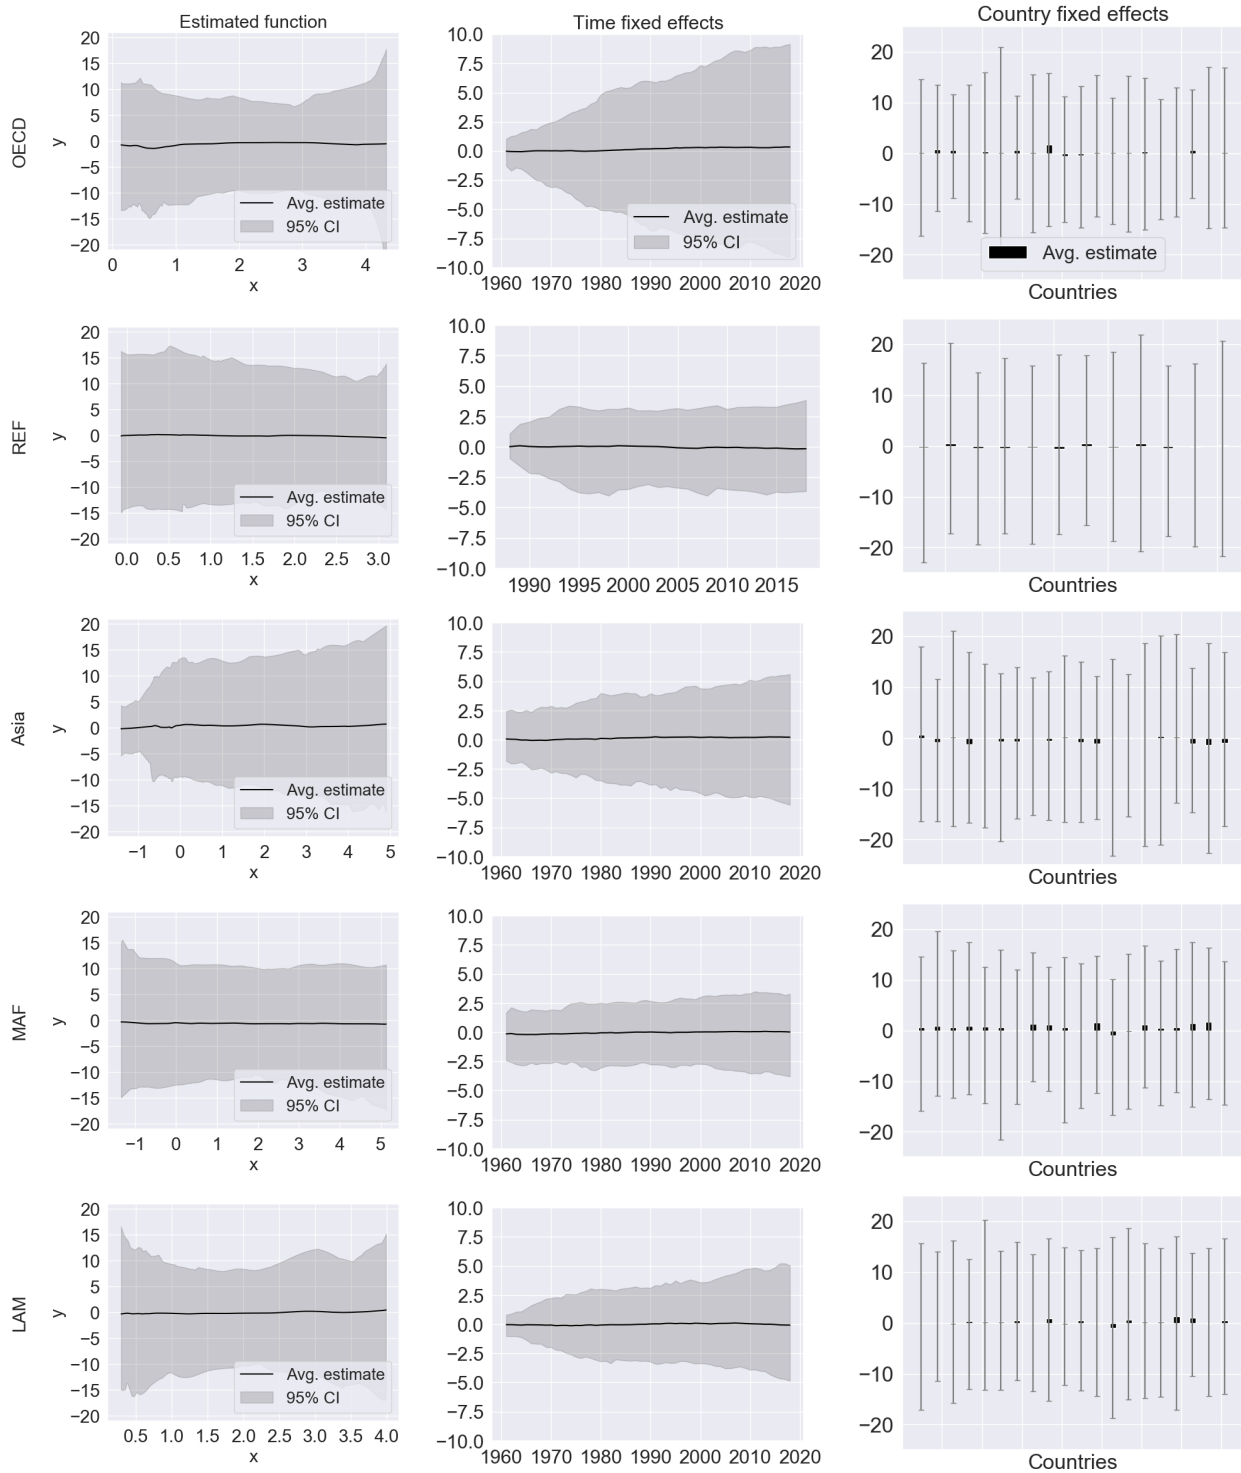
Note: In the country fixed effects plots, gray error bars indicate 95% confidence intervals.

Figure A.4: Monte Carlo spuriousness results for the regional formulation of the dynamic model
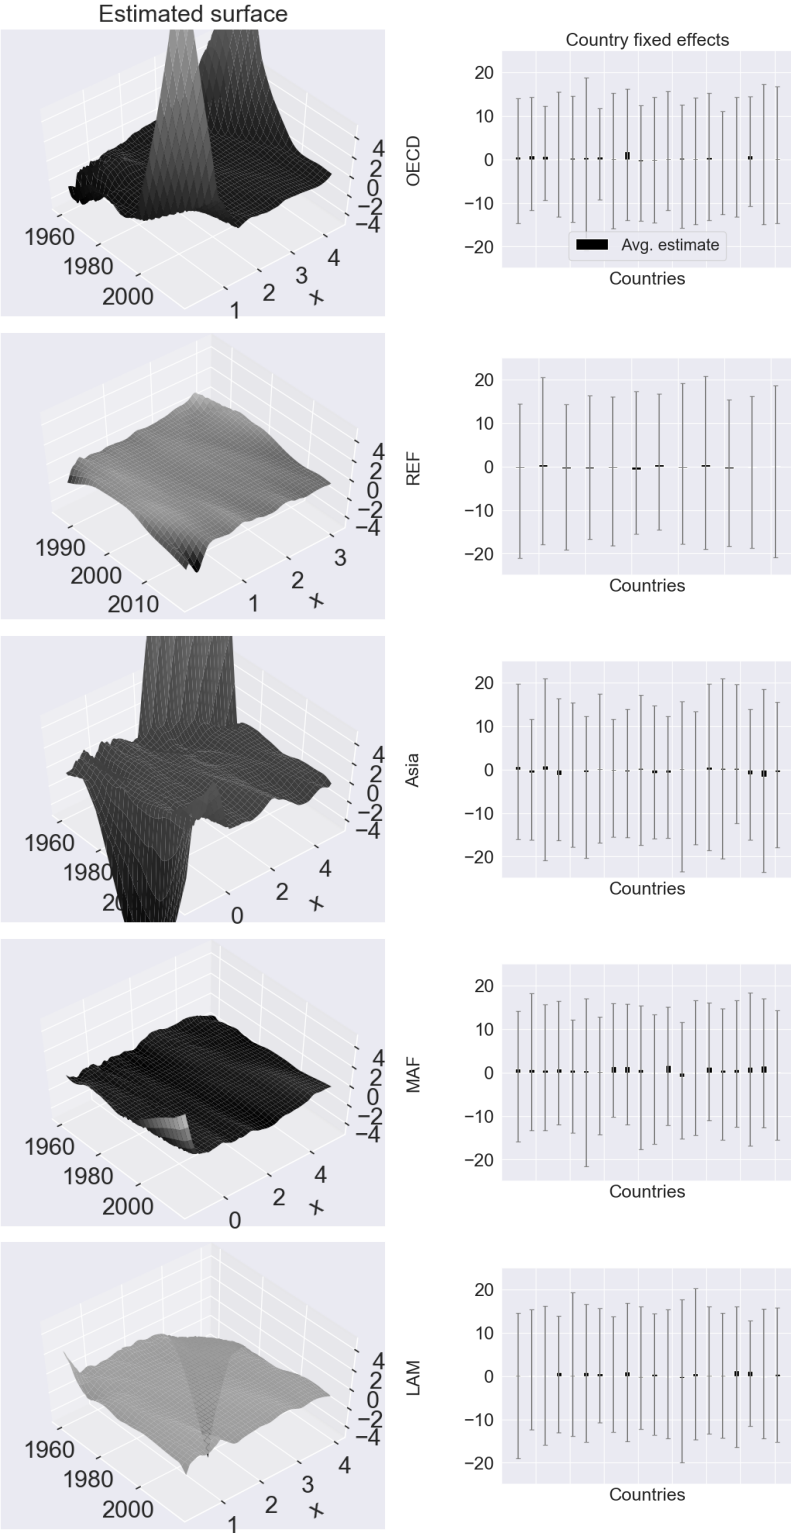
Note: In the country fixed effects plots, gray error bars indicate 95% confidence intervals.

Figure A.5: Estimation results for the benchmark models and consumption-based emissions
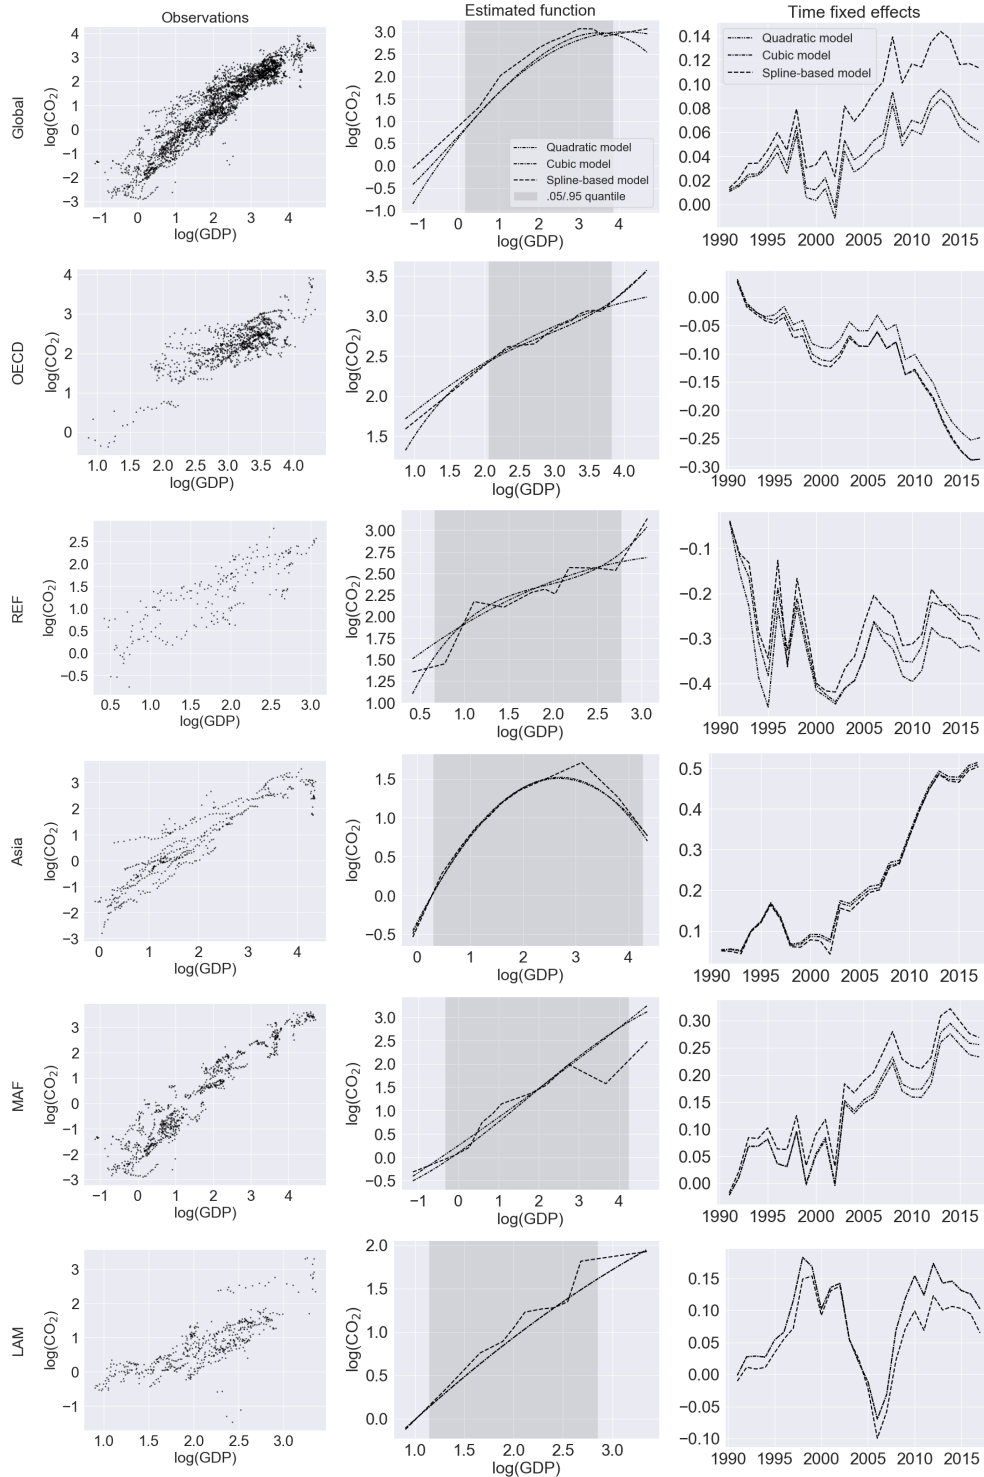
Note: Estimation results in the first row of plots are based on global model formulations where every country is used for estimation. Remaining rows show estimation results from region-wise estimation of the benchmark models.

Figure A.6: Estimation results for the dynamic model and territorial emissions based on a restricted sample that coincides with the one for consumption-based emissions

Note: Black coloring indicates area between the .05 and .95 income quantiles. Dark gray coloring indicates area with data support. Light gray coloring indicates area with no data support. The estimated surface in the top left panel is based on the global formulation of the dynamic neural network model. Remaining surfaces are based on the regional formulation of the dynamic neural network model.

A.2 Tables

Table A.1: Macro-region definitions

OECD

Albania[1,2]	Australia[1,2]	Austria[1,2]
Belgium[1,2]	Bosnia and Herzegovina[1]	Bulgaria[1,2]
Canada[1,2]	Croatia[1,2]	Cyprus[1,2]
Czech Republic[1,2]	Denmark[1,2]	Estonia[1,2]
Finland[1,2]	France[1,2]	Germany[1,2]
Greece[1,2]	Hungary[1,2]	Iceland[1]
Ireland[1,2]	Italy[1,2]	Japan[1,2]
Latvia[1,2]	Lithuania[1,2]	Luxembourg[1,2]
Malta[1,2]	Montenegro[1]	Netherlands[1,2]
New Zealand[1,2]	Norway[1,2]	Poland[1,2]
Portugal[1,2]	Puerto Rico	Romania[1,2]
Serbia	Slovakia[1,2]	Slovenia[1,2]
Spain[1,2]	Sweden[1,2]	Switzerland[1,2]
North Macedonia[1]	Turkey[1,2]	United Kingdom[1,2]
United States[1,2]		

REF

Armenia[1,2]	Azerbaijan[1,2]	Belarus[1,2]
Georgia[1,2]	Kazakhstan[1,2]	Kyrgyzstan[1,2]
Moldova[1]	Russia[1,2]	Tajikistan[1]
Turkmenistan[1]	Ukraine[1,2]	Uzbekistan[1]
Kosovo[1]		

Asia

Afghanistan[1]	Bangladesh[1,2]	Bhutan[1]
Brunei Darussalam[1,2]	Cambodia[1,2]	China[1,2]
South Korea[1,2]	Fiji[1]	French Polynesia
India[1,2]	Indonesia[1,2]	Laos[1,2]
Malaysia[1,2]	Maldives[1]	Micronesia[1]
Mongolia[1,2]	Myanmar[1]	Nepal[1,2]
New Caledonia	Pakistan[1,2]	Papua New Guinea[1]
Philippines[1,2]	Samoa[1]	Singapore[1,2]
Solomon Islands[1]	Sri Lanka[1,2]	Taiwan
Thailand[1,2]	Timor-Leste[1]	Vanuatu[1]
Vietnam[1,2]	Tuvalu[1]	Macao[1]
Marshall Islands[1]	Palau[1]	Hong Kong[1,2]
Tonga[1]	Kiribati[1]	

MAF

Algeria[1]	Angola[1]	Bahrain[1,2]
Benin[1,2]	Botswana[1,2]	Burkina Faso[1,2]
Burundi[1]	Cameroon[1,2]	Cabo Verde[1]
Central African Republic[1]	Chad[1]	Comoros[1]

Congo (DRC)[1]	Congo (RDC)[1]	Côte d'Ivoire[1,2]
Djibouti	Egypt[1,2]	Equatorial Guinea[1]
Eritrea[1]	Ethiopia[1,2]	Gabon[1]
Gambia[1]	Ghana[1,2]	Guinea[1,2]
Guinea-Bissau[1]	Iran[1,2]	Iraq[1]
Israel[1,2]	Jordan[1,2]	Kenya[1,2]
Kuwait[1,2]	Lebanon[1]	Lesotho[1]
Liberia[1]	Libya[1]	Madagascar
Malawi[1,2]	Mali[1,2]	Mauritania[1]
Mauritius[1,2]	Morocco[1,2]	Mozambique[1,2]
Namibia[1,2]	Niger[1]	Nigeria[1,2]
Palestine[1]	Oman[1,2]	Qatar[1,2]
Rwanda[1,2]	Saudi Arabia[1,2]	Senegal[1,2]
Sierra Leone[1]	Somalia	South Africa[1,2]
Sudan[1]	Eswatini[1]	Syria
Togo[1,2]	Tunisia[1,2]	Uganda[1,2]
United Arab Emirates[1,2]	Tanzania[1,2]	Yemen[1]
Zambia[1,2]	Zimbabwe[1,2]	Seychelles[1]
Sao Tome and Principe[1]		

LAM

Argentina[1,2]	Aruba[1]	Bahamas[1]
Barbados[1]	Belize[1]	Bolivia[1,2]
Brazil[1,2]	Chile[1,2]	Colombia[1,2]
Costa Rica[1,2]	Cuba	Dominican Republic[1,2]
Ecudaor[1,2]	El Salvador[1,2]	Grenada[1]
Guatemala[1,2]	Guyana[1]	Haiti[1]
Honduras[1,2]	Jamaica[1,2]	Mexico[1,2]
Nicaragua[1,2]	Panama[1,2]	Paraguay[1,2]
Peru[1,2]	Suriname[1]	Trinidad and Tobago[1,2]
Uruguay[1,2]	Venezuela[1,2]	Saint Vincent[1]
Curaçao[1]	Saint Lucia[1]	Antigua and Barbuda[1]
Dominica[1]	Bermuda[1]	

Note: [1]Country is in the CO_2 panel. [2]Country is in the CO_2^C and $CO_2^\star$ panel. Region definitions are from the Shared Socioeconomic Pathways (SSPs; Riahi et al., 2017): "OECD" is OECD90 and EU member states and candidates; "REF" is reforming economies Eastern Europe and the former Soviet Union; "Asia" is Asian countries excluding the Middle East, Japan and former Soviet Union states; "MAF" is the Middle East and Africa; and "LAM" is Latin America and the Caribbean. Countries in red are in the SSP database but not in our data set. Countries in green are in our data set but not in the SSP database.

Table A.2: Descriptive statistics: mean

	Mean						
	1960	1970	1980	1990	2000	2010	2017
Global							
GDP	3.70	5.49	6.78	7.69	9.01	11.46	13.50
CO_2	2.51	3.51	3.68	4.19	4.01	4.66	4.64
$CO_2^\star$	2.58	3.62	3.80	4.34	4.17	4.90	4.90
CO_2^C	NA	NA	NA	4.35	4.19	4.89	4.89
OECD							
GDP	10.84	16.09	20.47	24.76	29.45	32.62	35.99
CO_2	8.22	12.08	12.38	12.24	12.31	11.22	10.18
$CO_2^\star$	8.22	12.08	12.39	12.25	12.36	11.25	10.21
CO_2^C	NA	NA	NA	13.02	13.52	12.69	11.51
REF							
GDP	NA	NA	NA	10.29	6.52	11.16	12.30
CO_2	NA	NA	NA	14.06	7.69	8.62	8.44
$CO_2^\star$	NA	NA	NA	15.09	8.26	9.57	9.51
CO_2^C	NA	NA	NA	12.59	6.60	7.87	8.29
Asia							
GDP	0.78	1.00	1.37	2.14	3.52	6.54	9.51
CO_2	0.68	0.65	1.00	1.39	1.79	3.37	3.82
$CO_2^\star$	0.69	0.66	1.02	1.41	1.82	3.45	3.91
CO_2^C	NA	NA	NA	1.44	1.74	3.13	3.59
MAF							
GDP	3.23	5.61	6.84	5.06	5.55	6.86	7.21
CO_2	0.89	1.44	2.09	1.69	2.11	2.46	2.46
$CO_2^\star$	1.10	1.71	2.41	1.91	2.45	2.88	2.87
CO_2^C	NA	NA	NA	1.68	1.90	2.72	2.73
LAM							
GDP	5.03	6.54	9.13	8.53	9.84	11.86	12.44
CO_2	1.32	1.76	2.40	2.24	2.58	2.88	2.73
$CO_2^\star$	1.34	1.78	2.42	2.27	2.60	2.92	2.76
CO_2^C	NA	NA	NA	2.26	2.49	3.05	3.00

Note: "CO_2^C" are consumption-based CO_2 emissions; "$CO_2^\star$" are CO_2 emissions based on a restricted sample that coincides with the one for consumption-based emissions; and GDP is from the CO_2 panel. For given region r and time period t, mean values are calculated as the per capita values for that region: Mean $\equiv$ $\sum_{i \in I_r} \frac{x_{it}}{POP_{it}} \frac{POP_{it}}{POP_t^r} = \sum_{i \in I_r} \frac{x_{it}}{POP_t^r}$, where $I_r \subseteq \{1, 2, \ldots, N\}$ is the set of indices of countries belonging to region r, N is the total number of countries, POP_{it} is the population size, $POP_t^r \equiv \sum_{i \in I_r} POP_{it}$, and $x_{it} \in \{GDP_{it}, CO_{2it}, CO_2^\star{}_{it}, CO_2^C{}_{it}\}$.

Table A.3: Descriptive statistics: standard deviation

				Standard deviation			
	1960	1970	1980	1990	2000	2010	2017
Global							
GDP	4.71	7.24	9.43	9.79	11.35	11.67	12.39
CO_2	4.37	5.81	5.58	5.65	5.20	4.75	4.42
$CO_2{}^\star$	4.43	5.89	5.66	5.74	5.29	4.80	4.45
$CO_2{}^C$	NA	NA	NA	5.72	5.55	5.00	4.59
OECD							
GDP	4.16	4.09	5.33	7.88	10.21	10.03	10.51
CO_2	6.13	6.60	6.22	5.79	6.23	5.29	4.63
$CO_2{}^\star$	6.13	6.60	6.22	5.80	6.22	5.28	4.63
$CO_2{}^C$	NA	NA	NA	5.30	6.13	5.19	4.62
REF							
GDP	NA	NA	NA	3.26	2.62	4.62	5.14
CO_2	NA	NA	NA	4.36	2.93	3.86	4.30
$CO_2{}^\star$	NA	NA	NA	3.26	2.61	3.26	3.70
$CO_2{}^C$	NA	NA	NA	4.92	2.24	2.29	2.64
Asia							
GDP	0.51	0.78	1.74	2.28	3.25	4.78	6.11
CO_2	0.46	0.45	0.82	1.08	1.41	2.69	2.85
$CO_2{}^\star$	0.46	0.45	0.82	1.07	1.40	2.68	2.83
$CO_2{}^C$	NA	NA	NA	1.37	1.80	2.47	2.64
MAF							
GDP	2.23	9.60	14.06	7.95	9.16	9.77	10.25
CO_2	1.55	2.55	3.63	2.98	4.08	4.48	4.53
$CO_2{}^\star$	1.71	2.87	4.07	3.31	4.58	5.04	5.11
$CO_2{}^C$	NA	NA	NA	2.84	3.27	4.65	4.79
LAM							
GDP	2.46	2.89	2.90	2.55	2.99	3.36	3.27
CO_2	1.33	1.44	1.82	1.56	1.75	2.21	1.85
$CO_2{}^\star$	1.33	1.39	1.58	1.54	1.62	2.15	1.77
$CO_2{}^C$	NA	NA	NA	1.22	1.20	1.45	1.58

Note: "$CO_2{}^C$" are consumption-based CO_2 emissions; "$CO_2{}^\star$" are CO_2 emissions based on a restricted sample that coincides with the one for consumption-based emissions; and GDP is from the CO_2 panel. For given region r and time period t, standard deviations are calculated using the following relation: Standard deviation $\equiv \sqrt{\sum_{i \in I_r} \left(\frac{x_{it}}{POP_{it}} - \text{Mean} \right)^2 \frac{POP_{it}}{POP_t^r}}$, where Mean is from Table A.2, $I_r \subseteq \{1, 2, \dots, N\}$ is the set of indices of countries belonging to region r, N is the total number of countries, POP_{it} is the population size, $POP_t^r \equiv \sum_{i \in I_r} POP_{it}$, and $x_{it} \in \{\text{GDP}_{it}, CO_{2it}, CO_2{}^\star{}_{it}, CO_2{}^C{}_{it}\}$.

Table A.4: Neural network architectures

Network architecture	# parameters in static model			# parameters in dynamic model		
	Global $(R=1)$	Regional $(R=5)$	National $(R=186)$	Global $(R=1)$	Regional $(R=5)$	National $(R=186)$
(2)	6	14	376	8	16	378
(4)	12	28	752	16	32	756
(8)	24	56	1504	32	64	1512
(16)	48	112	3008	64	128	3024
(32)	96	224	6016	128	256	6048
(2,2)	12	20	382	14	22	384
(4,2)	20	28	390	24	32	394
(4,4)	32	48	772	36	52	776
(8,2)	36	44	406	44	52	414
(8,4)	56	72	796	64	80	804
(8,8)	96	128	1576	104	136	1584
(16,2)	68	76	438	84	92	454
(16,4)	104	120	844	120	136	860
(16,8)	176	208	1656	192	224	1672
(16,16)	320	384	3280	336	400	3296
(32,2)	132	140	502	164	172	534
(32,4)	200	216	940	232	248	972
(32,8)	336	368	1816	368	400	1848
(32,16)	608	672	3568	640	704	3600
(32,32)	1152	1280	7072	1184	1312	7104
(2,2,2)	18	26	388	20	28	390
(4,2,2)	26	34	396	30	38	400
(4,4,2)	40	48	410	44	52	414
(4,4,4)	52	68	792	56	72	796
(8,2,2)	42	50	412	50	58	420
(8,4,2)	64	72	434	72	80	442
(8,4,4)	76	92	816	84	100	824
(8,8,2)	108	116	478	116	124	486
(8,8,4)	128	144	868	136	152	876
(8,8,8)	168	200	1648	176	208	1656
(16,2,2)	74	82	444	90	98	460
(16,4,2)	112	120	482	128	136	498
(16,4,4)	124	140	864	140	156	880
(16,8,2)	188	196	558	204	212	574
(16,8,4)	208	224	948	224	240	964
(16,8,8)	248	280	1728	264	296	1744
(16,16,2)	340	348	710	356	364	726
(16,16,4)	376	392	1116	392	408	1132
(16,16,8)	448	480	1928	464	496	1944
(16,16,16)	592	656	3552	608	672	3568
(32,2,2)	138	146	508	170	178	540

(32,4,2)	208	216	578	240	248	610
(32,4,4)	220	236	960	252	268	992
(32,8,2)	348	356	718	380	388	750
(32,8,4)	368	384	1108	400	416	1140
(32,8,8)	408	440	1888	440	472	1920
(32,16,2)	628	636	998	660	668	1030
(32,16,4)	664	680	1404	696	712	1436
(32,16,8)	736	768	2216	768	800	2248
(32,16,16)	880	944	3840	912	976	3872
(32,32,2)	1188	1196	1558	1220	1228	1590
(32,32,4)	1256	1272	1996	1288	1304	2028
(32,32,8)	1392	1424	2872	1424	1456	2904
(32,32,16)	1664	1728	4624	1696	1760	4656
(32,32,32)	2208	2336	8128	2240	2368	8160

Note: "(a,b,c)" indicates a neural network architecture with three hidden layers containing a units in the first layer, b in the second, and c in the third; "# parameters" is the number of free model parameters excluding fixed effects.

Table A.5: Adam algorithm

Require: Learning rate ϵ (0.001)
Require: Exponential decay rates $\rho_1, \rho_2 \in [0, 1)$ for moment estimates (0.9, 0.999)
Require: Small constant δ used for numerical stabilization (10^{-8})
Require: Initial parameters θ
  Initialize first moment vector: $m = 0$
  Initialize second moment vector: $v = 0$
  Initialize time step: $t = 0$
  while stopping criterion not met **do**
      $t \leftarrow t + 1$
      Compute gradient: $g \leftarrow \nabla_\theta J(\theta)$
      Update biased first moment estimate: $m \leftarrow \rho_1 m + (1 - \rho_1)g$
      Update biased second moment estimate: $v \leftarrow \rho_2 v + (1 - \rho_2)g \odot g$
      Correct bias in first moment: $\hat{m} \leftarrow m/(1 - \rho_1^t)$
      Correct bias in second moment: $\hat{v} \leftarrow v/(1 - \rho_2^t)$
      compute update: $\Delta\theta = -\epsilon \cdot \hat{m}/(\sqrt{\hat{v}} + \delta)$
      Apply update: $\theta \leftarrow \theta + \Delta\theta$
  end while

Source: Adapted from Kingma and Ba (2014). Numbers in parentheses are suggested defaults.

The Adam algorithm is an extension to gradient descent that individually adapts the learning rate for all parameters in two ways. First, the algorithm scales the learning rate inversely proportional to the square root of an exponentially decaying average of past squared values of the gradient (second moment estimate). In this way, the learning rate for parameters with small partial derivatives of the loss function decreases less rapidly than for parameters with large partial derivatives. This implies greater progress in more gently sloped regions of the parameter space. Second, to speed up optimization, especially in face of pathological curvature, Adam incorporates so-called *momentum* by also scaling the learning rate proportionally to an exponentially decaying average of past values of the gradient (first moment estimate). Finally, Adam includes bias corrections of the moment estimates to account for initialization at the origin.

Table A.6: Optimal neural network architectures for the static model in the Monte Carlo misspecification experiment

	Function	
	Linear-quadratic	Linear-cubic
Global model formulation		
Neural network architecture	(2,2)	(2,2,2)
# parameters (excl. fixed effects)	12	18
Regional model formulation		
Neural network architecture	(4,4,4)	(4,4)
# parameters (excl. fixed effects)	52	32

Note: "(a,b,c)" indicates a neural network architecture with three hidden layers containing a units in the first layer, b in the second, and c in the third.

Table A.7: Optimal neural network architectures for the Monte Carlo spuriousness experiment

	Formulation	
	Global	Regional
Static model specification		
Neural network architecture	(4,2)	(4,4)
# parameters (excl. fixed effects)	20	48
Dynamic model specification		
Neural network architecture	(4,2)	(16,8)
# parameters (excl. fixed effects)	24	224

Note: "(a,b,c)" indicates a neural network architecture with three hidden layers containing a units in the first layer, b in the second, and c in the third.

Research Papers

2021-09:   Leopoldo Catania, Alessandra Luati and Pierluigi Vallarino: Economic vulnerability is state dependent

2021-10:   Søren Johansen and Anders Ryghn Swensen: Adjustment coefficients and exact rational expectations in cointegrated vector autoregressive models

2021-11:   Bent Jesper Christensen, Mads Markvart Kjær and Bezirgen Veliyev: The incremental information in the yield curve about future interest rate risk

2021-12:   Mikkel Bennedsen, Asger Lunde, Neil Shephard and Almut E. D. Veraart: Inference and forecasting for continuous-time integer-valued trawl processes and their use in financial economics

2021-13:   Anthony D. Hall, Annastiina Silvennoinen and Timo Teräsvirta: Four Australian Banks and the Multivariate Time-Varying Smooth Transition Correlation GARCH model

2021-14:   Ulrich Hounyo and Kajal Lahiri: Estimating the Variance of a Combined Forecast: Bootstrap-Based Approach

2021-15:   Salman Huseynov: Long and short memory in dynamic term structure models

2022-01:   Jian Kang, Johan Stax Jakobsen, Annastiina Silvennoinen, Timo Teräsvirta and Glen Wade: A parsimonious test of constancy of a positive definite correlation matrix in a multivariate time-varying GARCH model

2022-02:   Javier Haulde and Morten Ørregaard Nielsen: Fractional integration and cointegration

2022-03:   Yue Xu: Spillovers of Senior Mutual Fund Managers' Capital Raising Ability

2022-04:   Morten Ørregaard Nielsen, Wonk-ki Seo and Dakyung Seong: Inference on the dimension of the nonstationary subspace in functional time series

2022-05:   Kristoffer Pons Bertelsen: The Prior Adaptive Group Lasso and the Factor Zoo

2022-06:   Ole Linnemann Nielsen and Anders Merrild Posselt: Betting on mean reversion in the VIX? Evidence from ETP flows

2022-07:   Javier Hualde and Morten Ørregaard Nielsen: Truncated sum-of-squares estimation of fractional time series models with generalized power law trend

2022-08:   James MacKinnon and Morten Ørregaard Nielsen: Cluster-Robust Inference: A Guide to Empirical Practice

2022-09:   Mikkel Bennedsen, Eric Hillebrand and Sebastian Jensen: A Neural Network Approach to the Environmental Kuznets Curve