HANDBOOK ON ENERGY AND CLIMATE CHANGE

Handbook on Energy and Climate Change

Edited by

Roger Fouquet

Principal Research Fellow, Grantham Research Institute on Climate Change and the Environment, London School of Economics, UK

Edward Elgar Cheltenham, UK • Northampton, MA, USA © Roger Fouquet 2013

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical or photocopying, recording, or otherwise without the prior permission of the publisher.

Published by Edward Elgar Publishing Limited The Lypiatts 15 Lansdown Road Cheltenham Glos GL50 2JA UK

Edward Elgar Publishing, Inc. William Pratt House 9 Dewey Court Northampton Massachusetts 01060 USA

A catalogue record for this book is available from the British Library

Library of Congress Control Number: 2012943745

This book is available electronically in the ElgarOnline.com Economics Subject Collection, E-ISBN 978 0 85793 369 0



ISBN 978 0 85793 368 3 (cased)

Typeset by Servis Filmsetting Ltd, Stockport, Cheshire Printed by MPG PRINTGROUP, UK

1 INTRODUCTION

We identify anthropogenic contributions to atmospheric CO_2 measured at Mauna Loa using the statistical automatic model selection algorithm *Autometrics*. Estimating the determinants of atmospheric CO_2 is traditionally a challenge due to the complex systems of data involved. CO_2 is a highly autocorrelated, non-stationary time series, and globally there exist a large number of potential carbon sources and sinks. There is mixed evidence in the literature on anthropogenic contributions to atmospheric CO_2 : the long-term trend is widely attributed to human factors, while the main seasonal fluctuations are thought to be driven by the biosphere. However, the statistical measures applied are often somewhat unsatisfactory due to the complexities of dealing with large numbers of variables.

Over the long run of geological time, evidence of repeated glaciations, and of coal and oil deposits from extinct tropical forests, reveals that atmospheric CO_2 has varied greatly, and manifestly without any anthropogenic influence, including very low levels and levels as high as 1000 parts per million (ppm): see, for example, Hoffman and Schrag (2000); Hendry (2011) provides a summary. In the more recent half million years of 'ice ages', natural fluctuations include highs and lows of 300 and 180 ppm from Antarctic ice sheet drilling (see Juselius and Kaufmann, 2009). Finally, in the last 10 000–12 000 years, humanity has transformed planet Earth, replacing forests by agriculture and creating an industrial world (see, e.g., Ruddiman, 2005). Against the background of such movements, it is important to establish that recent levels of atmospheric CO_2 are not merely a natural event, but have an anthropogenic signature.¹

We introduce a new approach to modeling changes in atmospheric CO_2 using a model selection algorithm which allows for a larger number of potential variables than observations without *a priori* forcing any to be significant or to be excluded. Using this method, the main relevant explanatory variables are determined and their magnitudes estimated while irrelevant factors are dropped from the model.

The model controls for a number of natural carbon sources and sinks: vegetation measured by the Normalized Difference Vegetation Index (NDVI); temperature (measured as anomalies in the northern hemisphere); weather phenomena (measured through the Southern Oscillation Index); as well as accounting for dynamic transport by including seasonal interaction terms. This allows an estimate of the anthropogenic contribution to CO_2 as measured by industrial output indices and fossil fuel use for different geographical areas. The resulting estimates describe the direct effects on CO_2 growth and the proportional contribution of each factor. We find that vegetation, temperature and other natural factors alone cannot explain either the trend or all the variation in CO_2 growth. Industrial production components, driven by business cycles and shocks, are highly significant contributors.

Section 2 provides an overview of related literature, and Section 3 discusses model

selection, impulse indicator saturation (IIS) – which we use to detect multiple breaks in the models – and the *Autometrics* algorithm. Section 4 describes the data used, Section 5 outlines the estimation procedures and Section 6 reports the main results. Section 7 concludes.

2 LITERATURE REVIEW

There is a plethora of literature on atmospheric CO_2 and its link to anthropogenic factors. Key aspects in the literature are finding an appropriate measure for anthropogenic activity and sufficient controls for natural effects such as vegetation and oceanic absorption. Atmospheric CO_2 has been measured consistently and regularly since 1958, mostly due to the effort of Charles Keeling, who initiated and supported the measurement at Mauna Loa, Hawaii and later other measurement stations (see Sundquist and Keeling, 2009). This led to the now well-established and often cited 'Keeling curve', showing the increasing trend and highly seasonal pattern in atmospheric CO_2 (Figure 12.1). We are primarily interested in identifying the anthropogenic contribution to atmospheric CO_2 , so the following section reviews existing evidence and factors that need to be included in models.

Sources and Sinks

Anthropogenic sources

While atmospheric CO_2 has been consistently measured at multiple sites for a long time (see Keeling et al., 1976; Sundquist and Keeling, 2009), the choice of anthropogenic



Figure 12.1 Atmospheric CO₂ measured at Mauna Loa – the 'Keeling curve'

variables is less straightforward. Three variables are regularly used: estimates of fossil fuel emissions; population; and cement production. In most cases, variables are measured on an annual basis and interpolated to monthly frequencies.

A standard measure for estimating fossil fuel emissions is the product of the amount of fuel produced, the proportion of the fuel that is oxidized, and the fuel carbon content (Marland and Rotty, 1984). Variations of these data are used in Erickson et al. (2008), Jones and Cox (2005), Randerson et al. (1997), and Nevison et al. (2008). Recent data using this methodology are available at an annual frequency from Marland et al. (2011). In contrast, Hofman et al. (2009) as well as Newell and Marcus (1987) focus on population as a measure of human industrial output. Granados et al. (2012) extend the model of population by including a measure of global GDP. Cement production is a further major component of anthropogenic emissions. CO_2 emissions in production through limestone calcination, kiln operation and power generation are estimated to make up approximately 5 percent of anthropogenic emissions (Worrell et al., 2001).

These measures provide a good starting point to capture anthropogenic emissions in the long run, but due to the annual measurement frequency do not capture short-run fluctuations and seasonality. Population and GDP are too broad as measures to capture variability other than a trend. Measurement could be improved through supplementing overall fuel emissions by disaggregate individual fuel consumption. However, most importantly, a higher-frequency (monthly) anthropogenic measure is required to capture seasonal and short-term effects.

Terrestrial biosphere and transport

Aside from anthropogenic emissions, the terrestrial biosphere (vegetation) is a major factor in the carbon cycle's sources and sinks. When trying to model anthropogenic contributions to CO_2 it is therefore important to account for the most important factors in the natural world.

Atmospheric CO₂ falls and rises seasonally each year due to photosynthetic activity (during summer) and respiratory release (during winter) of CO₂ in the terrestrial biosphere (Keeling et al., 1996; Buermann et al., 2007). The intensity of these effects depends on the length of the growing season, a fertilization effect (feedback to plant growth from increased CO₂), and shifts in seasonal patterns (Kaufmann et al., 2008). The literature on atmospheric CO₂ suggests a strong link between the earth's biosphere and the level of CO₂.

The Normalized Difference Vegetation Index (NDVI) provides a direct measure of photosynthetic activity. As Myeni et al. (1995) describe, chlorophyll found in plants absorbs visible light for photosynthesis and reflects near infra-red light. The more active a plant is (indicated by higher density of green leaves), the more visible radiation is absorbed and the more near infra-red is reflected. Thus the difference between the two measures increases with higher leaf density. Using satellite-based remote sensing, the intensity of reflected visible and infra-red light can be measured. Using the ratio of the difference between the two measures, the NDVI is defined by

$$NDVI = \frac{\rho_{NearIR} - \rho_{Visible}}{\rho_{Visible} + \rho_{NearIR}}$$
(12.1)

where $\rho_{Visible}$ and ρ_{NearIR} are measures for visible and near infra-red light respectively. NDVI is therefore an index ranging from -1 to +1, with values around 0 denoting non-vegetation objects, and values around > 0.7 indicating dense vegetation (Tucker et al., 2010). The NDVI provides a direct measure for vegetation activity with high values close to 1 during the growing season (summer in the northern hemisphere), and lower values closer to 0 during the less active season (winter in the northern hemisphere). Kaufmann et al. (2008) investigate the link between NDVI and atmospheric CO₂ using econometric methods and find evidence that NDVI values 'Granger-cause' CO₂. There is also evidence of a feedback effect of increased CO₂, leading to enhanced vegetation activity.

Due to transport airflows, the primary influence of vegetation on measured atmospheric CO_2 depends on the location of the measurement station. In the case of Mauna Loa, Hawaii, the seasonal variation due to the biosphere is driven by long-range transport from Eurasia during winter and short-range transport from North America during summer (see Buermann et al., 2007; Taguchi et al., 2003). Narrowing down the time frame, Levin et al. (1995) suggest that Eurasian airflows dominate from October to June while North American airflows are dominant from July to September.

In terms of long-term development of terrestrial vegetation, there has been a greening trend – an increase of the growing season in the northern hemisphere (Lucht et al., 2002). This trend, however, was interrupted by the eruption of Mount Pinatubo in June 1991, which led to a decline in vegetation during 1992–93. This poses the question whether volcanic influence on atmospheric CO₂ needs to be controlled for separately from the biosphere. Lucht et al. (2002) state that the main channel through which volcanoes affect atmospheric CO₂ is indirect through temperature, while Hofman et al. (2009) propose that the Pinatubo eruption enhanced photosynthesis through scattered sunlight. Overall, Gerlach (2011) finds the direct effect of volcanic activity on measured atmospheric CO₂ to be small, thus accounting for vegetation, and temperature is expected to be sufficient without considering volcanoes separately. However, sudden breaks due to large eruptions will be detected through our IIS procedure.

Oceanic absorption and El Niño

A large amount of CO_2 is transferred between the ocean and the atmosphere, where the ocean acts as both a source and a sink of atmospheric CO_2 . There are two main factors determining oceanic carbon absorption and release (Bacastow, 1976). First, generally ocean absorption of CO_2 is governed by the difference in partial pressure between the atmospheric and oceanic CO_2 :

$$\frac{dp}{dt} = -k \cdot (p - P) \tag{12.2}$$

where p is the partial pressure of atmospheric CO_2 , P is the partial pressure of oceanic CO_2 beneath a surface layer and k is a variable capturing layer thickness and wind. These values are affected primarily by temperature and wind speed. Temperature is a key factor in oceanic CO_2 transfer, as higher sea surface temperature reduces uptake and increases outgassing (Watson et al., 1995). Increased wind speed increases k, so leads to higher oceanic CO_2 uptake. The second effect is upwelling – dense cold water driven to the

surface releases CO_2 stored in the ocean. Absorption and upwelling play opposite roles; which dominates is debated and depends on the geographical region.

The atmospheric fluctuations of air pressure differences known as Southern Oscillation affect oceanic absorption through the two channels described above. Southern Oscillation describes the change of air pressure differences between Tahiti and Darwin, Australia (see Troup, 1965; Bacastow, 1976; Keeling and Revelle, 1985). It is measured as an index (SOI) from the Australian Bureau of Meteorology (2011), and defined as

$$SOI = 10 \cdot \frac{\Delta P_t - \Delta P_t}{\sigma_{\Delta P_t}}$$
(12.3)

where ΔP_t is the difference in the average of mean sea level pressure between Tahiti and Darwin for month t. $\overline{\Delta P_t}$ is the long-run average of ΔP_t and $\sigma_{\Delta P_t}$ is the long-run standard deviation of ΔP_t for the given month respectively. Negative values of the SOI are generally referred to as El Niño years, while positive values correspond to episodes of La Niña. However, the effect on oceanic absorption is not so clear cut. Episodes of La Niña (SOI > 0) are associated with increased wind speeds (increase in k in equation (12.2)), thus making uptake easier. Nevertheless, increased wind also increases upwelling, which leads to a release of oceanic CO₂. Bacastow et al. (1985) suggest that easier absorption outweighs upwelling during episodes of La Niña (SOI > 0), resulting in higher absorption of CO₂ by the ocean when the SOI is positive. In turn, this implies less absorption during El Niño (SOI < 0) years.

On the contrary, Francey et al. (1995) find the opposite – during La Niña years (SOI>0) oceanic absorption is relatively lower because of the large upwelling effect. Keeling and Revelle (1985) side with Francey on the theoretical model that upwelling should outweigh increased absorption, but empirically find that less atmospheric CO_2 is absorbed during El Niño episodes (SOI<0), which agrees with Bacastow's (1976) findings.

Another factor that is not often considered in the literature is CO_2 use by oceanic algae, as Ritschard (1992) mentions. Nevertheless, data on algae are limited as they are not covered by the NDVI satellite measures, and consequently are not considered in our study.

Looking at the bigger picture, most evidence suggests that the ocean has become a carbon sink for anthropogenic emissions (Christopher et al., 2007). As increased CO_2 emissions into the atmosphere increased atmospheric partial pressure of CO_2 , absorption by the ocean increased due to the pressure difference. However, the absolute magnitude of this effect is not known, and they estimate that approximately 48 percent of fossil fuel emissions are absorbed in the ocean; and atmospheric CO_2 would be approximately 55 ppm (parts per million) higher if there were no oceanic uptake. Orr et al. (2001) similarly find the ocean to be a net carbon sink but, as well as Nevison et al. (2008), suggest that most models overestimate the proportion of CO_2 emissions absorbed by the ocean.

In order to capture the key factors determining oceanic uptake when modeling atmospheric CO_2 , it is important to account for temperature (which crucially affects partial pressure) and control for Southern Oscillation, even though the effects of El Niño and La Niña are not fully understood.

Modeling Methodology

The models for atmospheric CO_2 can broadly be classed into two categories: atmospheric transport and statistical models. Data for both are often decomposed into a long-run trend, cycle and noise using Fourier decompositions or the Hodrick–Prescott (HP) filter.

Atmospheric transport describes spatial three-dimensional models with vertical levels based on solving the fundamental equations for conservation of mass, momentum and energy. For different latitude and longitude grid resolutions, these models simulate carbon emissions and global transport (for the set-up and methodology of these models see Hansen et al., 1983; Kawa et al., 2004). Within this group of models, Erickson et al. (2008), Nevison et al. (2008), Randerson et al. (1997) and Keeling et al. (1995) use annual emissions data to analyze anthropogenic effects.

Statistical approaches also vary in methodology. Thoning and Tans (1989), Keeling et al. (1976) and Enting (1987) use Fourier decompositions to study trend and seasonal cycle. Granados et al. (2012) and Granados et al. (2009) use cointegration and timeseries regression to study links between population, GDP and HP-filtered CO_2 growth. However, no actual measure of anthropogenic emissions is used in these studies. Hofman et al. (2009) use regression and graphical comparisons of carbon and population, while Jones and Cox (2005) regress growth rates of CO_2 on global emissions and cement production. Newell and Marcus (1987) look at the simple correlation between levels of CO_2 and global population.

Concerning modeling methods, atmospheric transport models and many statistical techniques are widely applied and well documented. However, they suffer from similar problems. Often inappropriate statistical techniques are applied without considering the time-series properties of the data. Correlation of time series alone is not an appropriate measure of dependence between them. The low frequency of measurement of emissions data is problematic and models are restricted by an initial choice of a small number of independent variables. Original data are rarely used: instead series are decomposed. This step is not necessary *a priori*, especially when explanatory variables that are seasonal themselves are available. Additionally, the regression analyses applied in many papers are not robust to outliers or structural breaks, do not always handle non-stationarity, and present few tests for mis-specification.

Summary of the Main Findings

The long-run trend in increasing CO_2 is clearly driven by anthropogenic factors, whereas the short-run seasonal fluctuations and changes in amplitude are mostly attributed to changes in the biosphere.

Long-term trend

The long-term trend in atmospheric CO₂ is fossil fuel induced. Pre-industrial levels of CO₂ are estimated to be around 260–280 ppm (see Wigley, 1983; Hofman et al., 2009), based on ice core, tree ring and oceanic data. Consistent and repeated measurement, starting with Keeling's work in 1958, has documented the rise in CO₂ to a current level of approximately 390 ppm measured at Mauna Loa. The rate of increase of CO₂ is proportional to combustion of fossil fuels (see Keeling, 1973; Keeling et al., 1976; Keeling et

al., 1995; Thoning and Tans, 1989). Using population as a proxy measure for emissions yields similar results (Hofman et al., 2009; Newell and Marcus, 1987; Granados et al., 2009; Granados et al., 2012).

Given the large departure from pre-industrial levels, Hansen et al. (2008) investigate the question of a target level CO_2 . They suggest that (at the time of their writing) the level of 385 ppm is too high to maintain climate conditions that current life has adapted to. Levels of 450 ppm in the Caenozoic era were associated with near ice-free conditions. Consequently, they propose a target level of at most 350 ppm.

Seasonal variation and amplitude

Seasonal fluctuations and changes in amplitude are mainly attributed to factors in the biosphere rather than industrial emissions. There are two effects described in the literature: one is the general pattern of seasonality; the second is an increase in the amplitude of this seasonality. In particular a perceived increase in the growing season is alleged to be the driving force behind increases in amplitude.

Many studies propose that the seasonal component of atmospheric CO_2 reflects the inter-annual uptake by plants. This is supported by the fact that the amplitude of this seasonality for a given season decreases towards the equator (Keeling et al., 1976). In particular, Enting (1987) argues that vegetation is sufficient to account for most of the inter-annual variation and that economic data do not show the required seasonality. While the peak to trough ratio measured at Mauna Loa was approximately 0.8 for the time period Enting investigates, he suggests that industrial emissions are not sufficient to cause this seasonal change. However, as is obvious from many economic time series, there is high seasonality in production and therefore in emissions.

The amplitude of the seasonal effect has been increasing over time. Keeling et al. (1996), Randerson et al. (1997), Kohlmaier et al. (1989), and Bacastow et al. (1985) characterize the increase as a result of a lengthening growing season with only a very small effect directly from fossil fuel emissions. The effect from anthropogenic emissions in these studies ranges from 0.01 to 0.2 percent on the change of amplitude. Additionally, Keeling et al. (1995) find that changes in the overall growth rate of CO_2 are driven by changes in vegetation and temperature rather than changes in industrial emissions.

A major issue with many of the above-mentioned studies is that anthropogenic emissions and production data are measured annually and therefore do not have the required frequency to be able to account for seasonal fluctuations. In a recent paper, Erickson et al. (2008) investigate this issue and find that economic data would suggest that the highest anthropogenic fluxes occur at the same time as the respiration phase of plants (winter in the northern hemisphere). Once models account for seasonality in fuel consumption, this will then lead to a diminished effect of seasonality from the biosphere.

Contribution of this Chapter

There are recurring problems with existing models of anthropogenic contributions to CO_2 . Climate and atmospheric carbon fluxes are complex systems; nevertheless, many of the models are restricted by *a priori* selections of explanatory variables. The data used to account for anthropogenic emissions are often measured at too low a frequency to capture any seasonality. The main series of CO_2 is often decomposed into cycles and

trends, something that is not necessary if the explanatory data are measured at a reasonable frequency. On the one hand, a significant number of the papers that approach the problem from a statistics or economics point of view do not sufficiently control for the biosphere or other natural factors. On the other hand, many models coming from a natural science background use statistical methods that are ill fitted given the time-series characteristics of the data. Modern econometric methods can provide an interdisciplinary solution to these problems.

To address these issues, we introduce an extended general-to-specific (*Gets*) modeling approach based on automatic model selection and the theory of reduction. This allows for a large number of candidate explanatory variables; in particular, models can be estimated with more explanatory variables than observations. It is therefore possible to include many lags to capture time dynamics as well as a wide range of controls for natural factors and industrial output measures. As our main measure of anthropogenic productivity and emissions is industrial production (measured at monthly intervals), the data are analyzed as a whole without requiring prior decomposition into trends and cycles. Models are robust to outliers and structural breaks, handle unit roots reasonably well, and provide a straightforward method of testing for mis-specification.

Overall, the literature indicates a clear necessity to control for the biosphere, temperature, El Niño effects, and long- (as well as short-) run anthropogenic measures. Intuitively, our approach is to utilize a large number of potential determinants controlling for the above-mentioned factors, and then use automatic model selection techniques to determine which forces are significant. Starting with a theory-based, but very broad, general unrestricted model (GUM), the initial system is reduced to a specific model. This is a comparatively agnostic and data-driven approach that imposes few restrictions on explanatory variables while being robust to sudden shifts (structural breaks).

3 METHODOLOGY

The carbon cycle, with many potential sinks and sources, is a complex system, which makes it near-impossible to correctly specify an appropriate model *a priori*. To model complex equations, we rely on general to specific (*Gets*) modeling approaches (see Campos et al., 2005). The unknown data-generating process (DGP) is the underlying structure that creates the data. Empirical modeling will always deal with a subset of variables of the DGP; thus an important factor is the local data-generating process (LDGP) – the generating process in the space of the variables under analysis: see Hendry (2009). The approach, therefore, is to construct a set of data based on broad theoretical assumptions, which nests the LDGP; then, within this set, reduce the model from its general form down to a specific representation. This is a two-step procedure. One: define a set of *N* variables that includes the LDGP as a sub-model. Two: starting with that general model as a good approximation of the overall properties of the data, reduce its complexity by removing insignificant variables, while checking that at each reduction the validity of the model is preserved. This is the basic framework of *Gets* modeling.

This section introduces theoretical concepts of model selection, their use in misspecification testing, followed by the introduction of impulse indicator saturation (IIS) and its generalized version. All these concepts are then united and applied through the automatic search algorithm *Autometrics*. The algorithm combines these features through automated selection based on *Gets* while handling more variables than observations with IIS for detecting breaks and outliers, and mis-specification testing.

Model Selection

The theory of reduction characterizes the operations implicitly applied to the DGP to obtain the local LDGP. Choosing to analyze a set of variables, denoted $\mathbf{y}_t, \mathbf{x}_{1,t}$, determines the properties of the LDGP, and hence of any models of \mathbf{y}_t given $\mathbf{x}_{1,t}$ (with appropriate lags, non-linear transforms thereof etc.). A congruent model is one that matches the empirical characteristics of the associated LDGP, evaluated by a range of mis-specification tests (see, e.g., Hendry, 1995, and the following section). A model is undominated if it encompasses, but is not encompassed by, all other sub-models (see, e.g., Mizon and Richard, 1986; Hendry and Richard, 1989; and Bontemps and Mizon, 2008).

Mis-specification Testing

Using a large number of variables with IIS (discussed in the next section) also provides a new view of model evaluation: to avoid mis-specification and non-constancy, start as general as possible within the theoretical framework, using all the available data unconstrained by N > T (where N is the number of variables and T the number of observations), retaining the theory-inspired variables and only selecting over the additional candidates.

Even so, this approach does not obviate the need to test the specification of the auxiliary hypotheses against the possibility that the errors are not independent, are heteroskedastic (non-constant variance) or non-normal, that the parameters are not constant, that there is unmodeled non-linearity, and that the conditioning variables are not independent of the errors. When $N \ll T$, the first five are easily tested in the initial general model; otherwise their validity can be checked only after a reduction to a feasible sub-model. Congruence is essential not only to ensure a well-specified final selection, but also for correctly calibrated decisions during selection based on Gaussian significance levels, which IIS will help ensure.

More Variables than Observations: N > T

The model selection approach introduced here allows for more variables than observations to be used in modeling (N > T). For *Autometrics* this was first introduced through impulse indicator saturation, and has recently been extended to the general case.

Impulse indicator saturation

The numbers, timings and magnitudes of breaks in models are usually unknown, and are obviously unknown for unknowingly omitted variables, so a 'portmanteau' approach is required that can detect location shifts at any point in the sample while also selecting over many candidate variables. To check the null of no outliers or location shifts in a



Figure 12.2 Impulse indicator saturation under the null of no break

model, impulse indicator saturation (IIS) creates a complete set of indicator variables $\{1_{\{j=t\}}\}\$ equal to unity when j = t and equal to zero otherwise for $j = 1, \ldots, T$, and includes these in the set of candidate regressors. Although this creates T variables when there are T observations, in the 'split-half' approach analyzed in Hendry et al. (2008), a regression first includes T/2 indicators. By dummying out the first half of the observations, estimates are based on the remaining data, so any observations in the first half that are discrepant will result in significant indicators. The location of the significant indicators is recorded, then the first T/2 indicators are replaced by the second T/2, and the procedure repeated. The two sets of significant indicators are finally added to the model for selection of the significant indicators. The distributional properties of IIS under the null are analyzed in Hendry et al. (2008), and extended by Johansen and Nielsen (2009) to both stationary and unit-root autoregressions.

Figure 12.2 illustrates the 'split-half' approach for $y_t \sim IN[\mu, \sigma_y^2]$ for T = 100 selecting indicators at a 1 percent significance level (denoted α). The three rows correspond to the three stages: the first half of the indicators, the second half, then the selected indicators combined. The three columns respectively report the indicators entered, the indicators finally retained in that model, and the fitted and actual values of the selected model. Initially, although many indicators are added, only one is retained. When those indicators are dropped and the second half entered (row 2), none is retained. Now the combined retained dummies are entered (here just one), and selection again retains it. Since $\alpha T = 1$, that is the average false null retention rate.

We next illustrate IIS for a location shift of magnitude λ over the last k observations:

$$y_t = \mu + \lambda \mathbf{1}_{\{t \ge T - k + 1\}} + \varepsilon_t \tag{12.4}$$





Figure 12.3 IIS stages for a location shift

where $\varepsilon_t \sim IN [0, \sigma_{\varepsilon}^2]$ and $\lambda \neq 0$. The optimal test in this setting would be a *t*-test for a break in (12.4) at T - k + 1 onwards, but requires precise knowledge of the location-shift timing, as well as knowing that it is the only break and is the same magnitude break thereafter. Figure 12.3 records the behavior of IIS for a mean shift in (12.4) of $10\sigma_{\varepsilon}$ occurring at 0.75T = 75. Initially, many indicators are retained (top row), as there is a considerable discrepancy between the first-half and second-half means. When those indicators are dropped and the second set entered, all those for the period after the location shift are now retained. Once the combined set is entered (despite the large number of dummies), selection again reverts to just those for the break period. Under the null of no outliers or breaks, any indicator that is significance can be transient, due to an unmodeled feature that occurs elsewhere in the data set. Thus there is an important difference between 'outlier detection' procedures and IIS.

While IIS is perhaps surprising initially, many well-known statistical procedures are variants of IIS. The Chow (1960) test corresponds to subsample IIS over T - k + 1 to T, but without selection, as Salkever (1976) showed, for testing parameter constancy using indicators. Recursive estimation is equivalent to using IIS over the future sample, and reducing the indicators one at a time. Johansen and Nielsen (2009) relate IIS to robust estimation, and show that under the null of no breaks, outliers or data contamination, the average cost of applying IIS is the loss of αT observations. Thus, at $\alpha = 0.01$, for T = 100 one observation is 'dummied out' by chance despite including 100 'irrelevant' impulse indicators in the search set and checking for location shifts and outliers at every data point. Retention of theory variables, and IIS can be generalized to multiple splits of unequal size. While IIS entails more candidate variables than observations as N + T > T, selection is feasible as *Autometrics* undertakes expanding as well as contracting block



Figure 12.4 IIS outcomes for four location shifts

searches (see next section). Non-linear model selection (including threshold models) is examined in Castle and Hendry (2011).

For a single location shift, Hendry and Santos (2010) show that the detection power is determined by the magnitude of the shift τ ; the length of the break interval $T - T^*$, which determines how many indicators need to be found; the error variance of the equation σ_n^2 ; and the significance level, α , where a normal-distribution critical value, c_{α} , is used by the IIS selection algorithm. Castle et al. (2012b) establish the ability of IIS to detect multiple location shifts and outliers, including breaks close to the start and end of the sample, as well as correcting for non-normality. Figure 12.4 shows the application of the general autometrics algorithm to a trending process with four breaks of varying magnitudes over $1, \ldots, 10; 40, \ldots, 45; 75, \ldots, 90;$ and $90, \ldots, 100$, to illustrate the ability of IIS to capture multiple breaks, at both the start and end of the sample.

General case of N > T

The idea of generalizing using more variables than observations from IIS to all forms of independent variables is introduced by Hendry and Krolzig (2005) as well as Hendry and Johansen (2013). Suppose there are N total regressors partitioned into J blocks of n_j , where $N = \sum_{j=1}^{J} n_j$ such that N > T and $n_j < T$ for all j. Consequently the total number of variables N exceeds the number of observations T but total variables can be partitioned into J blocks of n_j , applying Gets to each block retaining the selected variables and crossing the groups to mix variables. The next step is to use the union of selected variables from each block to form a new initial model and repeat the process until the final union of selected variables is sufficiently small. Autometrics implements a variant of this algorithm to handle the general case of N > T.

Autometrics

Autometrics (see Doornik, 2009) is the latest installment in the automated Gets methodology and is available in the OxMetrics software package. The algorithm is based on the following main components:

- 1. GUM: the general unrestricted model (GUM) is the starting point of the search. The GUM should be specified based on broad theoretical considerations to nest the LDGP.
- 2. Pre-search: prior to specific selection, a pre-search lag reduction is implemented to remove insignificant lags, speeding up selection procedures and reducing the fraction of irrelevant variables selected (denoted the gauge of the selection process). Pre-search is only applied if the number of variables does not exceed the number of observations (N < T).
- 3. Search paths: *Autometrics* uses a tree search to explore paths. Starting from the GUM, *Autometrics* removes the least significant variable as determined by the lowest absolute *t*-ratio. Each removal constitutes one branch of the tree. For every reduction, there is a unique sub-tree which is then followed; each removal is back-tested against the initial GUM using an *F*-test. If back-testing fails, no sub-nodes of this branch are considered (though different variants of this removal exist). Branches are followed until no further variable can be removed at the pre-specified level of significance α . If no further variable can be removed, the model is considered to be terminal.
- 4. Diagnostic testing: each terminal model is subjected to a range of diagnostic tests based on a separately chosen level of significance. These tests include tests for normality (based on skewness and kurtosis), heteroskedasticity (for constant variance using squares), the Chow test (for parameter constancy in different samples), and residual autocorrelation and autoregressive conditional heteroskedasticity. Parsimonious encompassing of the feasible general model by sub-models both ensures no significant loss of information during reductions, and maintains the null retention frequency of *Autometrics* close to α: see Doornik (2008). Both congruence and encompassing are checked by *Autometrics* when each terminal model is reached after path searches, and it backtracks to find a valid less-reduced earlier model on that path if any test fails. This repeated re-use of the original mis-specification tests as diagnostic checks on the validity of reductions does not affect their distributions (see Hendry and Krolzig, 2003).
- 5. Tiebreaker: as a result of the tree search, multiple valid terminal models can be found. The union of these terminal models is referred to as the terminal GUM. As a tiebreaker to select a unique model, the likelihood-based Schwarz (1978) information criterion (SIC) is used, although other methods are also applicable, and terminal models should be considered individually.

In simulation experiments, models are primarily evaluated based on three concepts: gauge, potency and the magnitudes of the estimated parameters' root mean-square errors (RMSEs) around the DGP values (see Doornik and Hendry, 2009). Gauge describes the retention of irrelevant variables when selecting (i.e. variables that are selected but do not feature in the DGP). Potency measures the average retention frequency of relevant variables (variables that are selected and feature in the DGP). Low gauge (close to zero) and high potency (close to 1) are preferred, as are small RMSEs.

The main calibration decision in the search algorithm is the choice of significance level α at which selection occurs. Selection continues until retained variables are significant at α , although it can be the case that variables in the final model are also retained at a level

above α if removal leads to diagnostic tests failing. α is approximately equal to the gauge of selection. Further, the choice of diagnostic tests and lag length selection for residual autocorrelation and autoregressive conditional heteroskedasticity need to be set.

In the general case of N > T and IIS, *Autometrics* groups variables into two categories: selected and not selected (Doornik, 2010). Not currently selected variables are split into sub-blocks and the algorithm proceeds by alternating between two steps: first, the expansion step – selection is run over not-selected sub-blocks to detect omitted variables. Second, the reduction step – a new selected set is found by running selection on the system augmented with the omitted variables found in step one. This is repeated until the dimensions of the terminal model are small enough and the algorithm converges, so the final model is unchanged by further searches for omitted variables.

Autometrics has been applied successfully in a range of fields: see, for example, Hendry and Mizon (2011) on US food expenditure, Bårdsen et al. (2010) on unemployment in Australia, and Castle et al. (2011) for a comparison with other selection methods.

Nevertheless, overall selections should be interpreted carefully. Successful identification of the underlying LDGP can be adversely affected by collinearity of the independent variables. Most simulations of *Autometrics* with large numbers of variables use orthogonal regressors, which makes selection easier. Furthermore, when N > T in the block selection algorithm of *Autometrics*, adding or dropping a variable from the initial GUM may change the block partitioning of variables, so the selection is not invariant to the initial specification.

The next section covers the data used to construct the GUM in an attempt to nest the LDGP for atmospheric CO_2 . In the section following, *Autometrics* is then used to determine the anthropogenic contributions to CO_2 .

4 DATA

CO₂

The atmospheric CO_2 data used here are taken from Keeling's measurements at Mauna Loa, available from Tans and Keeling (2011) (Scripps Institution of Oceanography). The time series of CO_2 in monthly averages runs from 1958:3 until 2011:7 at the time of writing. Simple inspection of the data shows that both the level (see Figure 12.1) and the annual change (see Figure 12.5) are increasing over time.

The seasonal fluctuations are apparent in the data, and as Buermann et al. (2007) poetically describe it, the regular seasonal cycle of CO_2 at Mauna Loa 'records the breathing of the Northern Hemisphere biosphere'. In economic terms the level of atmospheric CO_2 can plainly be described as a stock variable. The total level of CO_2 can be approximated by the integral of the netflow to the atmosphere:

$$\operatorname{CO}_{2_t} = \int \operatorname{Netflow} dt$$
 (12.5)

where Netflow = carbon sources – carbon sinks. Therefore any analysis of the impact of anthropogenic emissions should be based on the change in atmospheric CO_2 :



Annual percentage changes in Mauna Loa CO₂

Figure 12.5 Annual changes in atmospheric CO₂

$$\frac{dCO_2}{dt} = \text{Netflow}$$
(12.6)

Approximating this relationship from continous to discrete time:

$$\frac{d\mathrm{CO}_2}{dt} \approx \Delta \mathrm{CO}_{2,t} = \mathrm{CO}_{2,t} - \mathrm{CO}_{2,t-1} = \mathrm{Netflow}_t$$
(12.7)

The dependent variable modeled is therefore $\Delta CO_{2,l}$. Equations (12.5)–(12.7) suggest that any relationship between the change in atmospheric CO_2 and netflow should be modeled in levels rather than any non-linear transformation thereof. The following section identifies variables that make up the netflow, both anthropogenic as well as natural sources and sinks.

Terrestrial Biosphere

We use NDVI data to account for vegetation effects on CO_2 . Data are available for the NDVI from the Oak Ridge National Laboratory Distributed Archive Center (see Tucker et al., 2010), ranging from 1981:7 until 2002:12 at spatial resolutions of 0.25, 0.5 and 1.0 degrees latitude and longitude.

 CO_2 measured at Mauna Loa is driven by North American airflow during the summer and airflow from Eurasia during the winter. Therefore the NDVI data are split into two main regions: North America and Eurasia. Figure 12.6 shows the NDVI measure in detail for North America in August and January respectively. Using 1° spatial resolution, an algorithm then takes the average of every 3 × 3 observation grid on land within the two regions (excluding water, permanent ice and missing observations). To capture



Figure 12.6 NDVI at 1° North America in 1981: August (top) and January (bottom)

North America	Proportion of variance	Cumulative
Principal component 1	0.831	
Principal component 2	0.064	0.895
Principal component 3	0.042	0.937
Eurasia		
Principal component 1	0.735	
Principal component 2	0.108	0.843
Principal component 3	0.042	0.885

Table 12.1 Principal components for vegetation, 1981:7–2002:12

the main variation of vegetation, similar to the spatial analysis in Buermann et al. (2007), the North American region is defined by the rectangle ranging from $86^{\circ}N/167^{\circ}W$ to $14^{\circ}N/48^{\circ}W$, and the Eurasian region by the rectangles ranging from $75^{\circ}N/9^{\circ}E$ to $36^{\circ}N/51^{\circ}E$ and $76^{\circ}N/52^{\circ}E$ to $7^{\circ}N/358^{\circ}E$. This generates 198 time-series variables for North America and 567 for Eurasia.

Due to the nature of a common growing season in the northern hemisphere, the generated time series are highly collinear. Principal components (PCs) are used to reduce the number of variables, while retaining most of the variation in the data. Since this process captures the overall variation in vegetation, it should also reduce the problem of random noise due to cloud cover at the time of satellite measurement. Although PCs are just linear transformations of the original times series, they have two potential advantages. First, PCs are mutually orthogonal, so adding or eliminating any one PC has little effect on the coefficient estimates of others, making the model more robust. Second, linear combinations of 'small' effects can be statistically significant (so retained during model selection) when individual time series would not be: see Castle et al. (2012a) for a more detailed discussion. The contributions of individual variables can be disentangled if needed.

For the following analysis, the first three principal components are entered for both North America and Eurasia. Cumulatively they explain 93.6 percent of the variation in North America and 88.5 percent of Eurasian variation. Table 12.1 summarizes the principal components that account for variation in the biosphere, and Figure 12.7 shows the highly seasonal variation present in the biosphere as measured by PCs. The first principal component shows a higher amplitude in Eurasia compared to North America but the seasonal pattern is nearly identical. NDVI implicitly covers changes in land use since it is a measure of photosynthetic activity for a particular area. A forest that is cut down would result in a change of NDVI from around +1 to closer to zero for that particular region. However, once NDVI is calculated for large regions and reduced in dimensionality (by PCs), changes in land use would have to occur on a grand scale to be identified in the time series. The principal components of NDVI should, therefore, be interpreted primarily as the variation in plant activity of photosynthesis and respiratory release.



Figure 12.7 First principal components for Eurasian and North American NDVI

Oceanic Indicators: Temperature and Southern Oscillation

The measure of temperature used here is the anomaly in land and sea surface temperature for the northern hemisphere. The temperature anomaly measure is expected to capture the main factors of ocean CO_2 absorption and is available from the NASA Goddard Institute for Space Studies (GISS) (2011) Surface Temperature Analysis from 1880:1 to 2011:9. The data are measured as an index in 0.01 degrees Celsius of deviations from the 1951–80 base period (see Hansen and Lebedeff, 1987; Hansen et al., 2010 for the detailed measurement methodology). Land measures are taken from multiple stations and are combined and corrected for urban and other non-climatic factors. Sea surface temperature measures are restricted to ice-free regions. As Hansen et al. (2010) describe, temperature in the northern hemisphere has been increasing despite recent El Niño effects. Figure 12.8 shows temperature anomalies from 1958:3 to 2010:9.

The feedback effect of CO_2 is one of the main concerns in climate change. The level of atmospheric CO_2 feeds into temperature which, in turn, affects the rate of growth of CO_2 , particularly through oceanic absorption. There would be a potential problem of endogeneity if the level of CO_2 were modeled by temperature. However, lagged temperature measures should be predetermined for the growth of CO_2 .

To take account of weather phenomena through the Southern Oscillation we include the Southern Oscillation index (SOI). Data on the SOI are available from the Australian Bureau of Meteorology (2011) from 1876:1 until 2011:9. Figure 12.9 shows the SOI for 1958 until 2011, with a noticeably strong episode of El Niño in 1997–98 (SOI < 0).



Figure 12.8 Land and sea temperature anomalies for the northern hemisphere



Figure 12.9 Southern Oscillation index

Economic Indicators

Anthropogenic contributions to atmospheric CO_2 are normally approximated by lowfrequency economic indicators; for example annual GDP, population, an estimate of total CO_2 emissions or cement production (see Section 2). This works reasonably well when trying to explain the long-run dynamics of carbon. However, using only low-frequency annual measures does not allow for estimation of any effect of anthropogenic emissions on the seasonal variation of CO_2 . High-frequency (monthly) measures of anthropogenic output permit a richer analysis. Here we use a combination of multiple low-frequency (annual) and high-frequency (monthly) indicators. The annual data are included to provide a robustness check and potentially account for long-run growth. Monthly measures are included for short-run dynamics, which could explain the seasonal fluctuations as well as long-run growth. To capture atmospheric transport, variables are chosen to reflect North America as well as Europe/Asia.

High-frequency (monthly) measures

The main high-frequency indicators for anthropogenic contribution to CO₂ used here are monthly industrial production indices for multiple regions. Industrial production for North America is given by the US Industrial Production Index (2005 = 100) available from the Federal Reserve (2011). The data are not seasonally adjusted and range from 1919:1 until 2011:5. The index measures real output in the sectors covering manufacturing, electric and gas utilities, and mining, and thus accounts for a large share in business-cycle fluctuations. To cover Europe and Asia, industrial production indices for the UK, Germany, India and Japan are included. These measures function as a proxy for business-cycle fluctuations in the Eurasian region. Ideally Chinese and Russian production should also be included; however, no data are available on industrial production for both countries at the required frequency and time span.² UK industrial production is measured as an index (2005 = 100) of non-seasonally adjusted manufacturing activities. The data are obtained from the Office of National Statistics (2011) and range from 1968:1 to 2011:6. German and Japanese industrial production is measured by the industrial production index (2005 = 100) covering manufacturing, mining and electricity, gas and water supply. The series is available only in seasonally adjusted format from the OECD (2011) from 1960:1 until 2011:2, and for Germany after October 1990 the data account for the accession of the German Democratic Republic to West Germany. The Indian industrial production index (2005 = 100) covers manufacturing, mining and electricity (Ministry of Statistics, Government of India 2011) and ranges from 1981:4 to 2011:5. Table 12.2 and Figure 12.10 summarize the high-frequency measures.

The seasonal adjustment of German and Japanese industrial production is visible in their dampened seasonal cycles. The higher seasonal variation in UK industrial production likely stems from it covering primarily manufacturing, which is more volatile to business cycles and seasonal factors than mining and energy production included in the other indices.

Measure	Description	Range	Source
US industrial production	Index $2005 = 100$	1958:3–2011:5	US Federal Reserve
UK industrial production	Index $2005 = 100$	1968:1–2011:6	ONS
Germany industrial production	Index $2005 = 100$	1960:1-2011:2	OECD
India industrial production	Index $2005 = 100$	1981:4-2011:5	Govt. of India
Japan industrial production	Index $2005 = 100$	1960:1-2011:2	OECD

Table 12.2 High-frequency (monthly) variables



Figure 12.10 Industrial production indices, 2005 = 100

<i>Table 12.3</i>	Principal	components	for industrial	production,	1981:4-2011:2
-------------------	-----------	------------	----------------	-------------	---------------

IP	Proportion of variance	Cumulative
Principal component 1	0.809	
Principal component 2	0.119	0.928
Principal component 3	0.047	0.975

These industrial production series are highly collinear, so we again employ principal components to reduce the dimensionality and work with orthogonal variables, which improves robustness in selection as discussed above. Table 12.3 summarizes the first three industrial production components, which cumulatively explain approximately 97 percent of variation in production. Figure 12.11 compares US industrial production with the three components used in selection.

These are the first three anthropogenic high-frequency components included in the GUM. While industrial production reflects the intensity of economic activity associated with higher emissions, it does not account for changing emission intensity. More efficient processes could lead to an increase in industrial production without an equivalent increase in emissions. This is a crucial missing measure and difficult to control for: an attempt is made by including overall long-run emissions estimates in addition to production.

Low-frequency (annual) measures

Low-frequency anthropogenic measures are variables reported on an annual basis, and capture the long term of human-driven CO_2 emissions. These low-frequency variables include total estimated CO_2 emissions (in thousand metric tons of carbon) for North



Figure 12.11 Industrial production principal components

Table 12.4	Low-frequency	(annual)	variables
		\ /	

Measure	Description	Range	Source
N. America CO_2 emissions	000s tons carbon	1950–2007	US Dep. of Energy
W. Europe CO_2 emissions	000s tons carbon	1950–2008	US Dep. of Energy
E. Europe CO_2 emissions	000s tons carbon	1950–2008	US Dep. of Energy
Central Asia CO_2 emissions	000s tons carbon	1950–2008	US Dep. of Energy
Far East CO_2 emissions	000s tons carbon	1950–2008	US Dep. of Energy

America, Western Europe, Eastern Europe, Central Asia and the Far East.³ Emissions are estimated based on the burning of fossil fuels, cement manufacture and gas flaring by the US Department of Energy (see Marland et al., 2011) from 1950 to 2007. Table 12.4 summarizes the annual measures.

In order to usefully combine these annual variables with the monthly measures listed above, all annual variables are linearly interpolated to monthly observations. In the case of moving from annual to monthly observations, this method estimates a straight line over 12 months between each annual observation. While this may be restrictive, if the seasonal structure of the variables is not known there is no *a priori* reason to prefer a different interpolation algorithm. In any case, seasonal dummy variables and IIS can 'pick up' any systematic or large deviations. As before, the low-frequency measures are reduced in dimensionality by calculating their principal components. The first three components capture 99 percent of variation in the emissions series. Table 12.5 shows



Table 12.5 Low-frequency (annual) components

Figure 12.12 CO, emissions and principal components

explained variation of each component and Figure 12.12 displays interpolated CO_2 emissions as well as the three low-frequency PCs.

The components of the interpolated annual variables are included to account for potential low-frequency movements, the industrial production short-term indicators are expected to be sufficient to capture inter-annual dynamics. These are the first three low-frequency anthropogenic components included in the GUM.

5 ESTIMATION

Overview

The dependent variable ΔCO_2 is modeled as a finite autoregressive-distributed lag model (ADL) (see, e.g., Hendry, 1995):

$$\Delta \text{CO}_{2,t} = a(L)\Delta \text{CO}_{2,t-1} + \sum_{i=1}^{p} b_i(L) x_{i,t} + \mathbf{z}_t' \mathbf{\phi} + \mathbf{\varepsilon}_t$$
(12.8)



Figure 12.13 Dependent variable ΔCO_{2t} *over 1982:7–2002:12*

where L denotes the lag operator, and p is the number of explanatory variables x_i . The vector z consists of non-lagged deterministic terms such as a linear time trend, centered seasonal variables and impulse indicators. Let q denote the total number of explanatory variables that appear on the left-hand side of equation (12.8). Then the change in atmospheric CO₂ is modeled as a function of past values of the change in CO₂, current and past values of selected independent variables x_i , and deterministic components. Figure 12.13 graphs $\Delta CO_{2,t}$ for 1982:7–2002:12. The seasonal variation is so large that it is difficult to discern the slow but persistent growth.

There is a large number of potential explanatory variables x_i in modeling atmospheric CO₂. The model needs to account for all the above-mentioned anthropogenic and natural factors as well as their lag reactions, as CO₂ is a highly autocorrelated time series. Adding IIS quickly moves the general unrestricted model to a situation with more explanatory variables than observations. This used to be a major difficulty in modeling; however, as outlined in Section 3, it can now be handled by estimation in blocks using *Autometrics* to select variables to retain in the final model in the form of (12.8). The estimation procedure operates as follows: first the theory-motivated GUM is specified, then estimation in blocks following the *Autometrics* algorithm selects down to individual terminal models. The union of terminal models is captured by the final GUM. Formally, the selection for the final model can use the likelihood-based SIC, although, as each terminal model represents a valid representation, final model selection can be based on other theoretical considerations.

Formulation of the GUM

The dependent variable that is being modeled is ΔCO_2 (the change in atmospheric CO_2 measured at Mauna Loa). While *Autometrics* has been shown to be effective at recovering the data-generating process in large models, the algorithm is not perfect and is sensitive to initial specification. As a robustness check, therefore, we estimate multiple variations of initial sample specifications. The subsample GUMs 1 and 2 are selected



Figure 12.14 Partial-autocorrelation function for ΔCO ,

for different initial variables, although they always include the main variables of interest: short-term dynamics captured by monthly industrial production and the control variables for natural factors – temperature, SOI and vegetation. Non-anthropogenic emissions are captured through potential oceanic release of CO_2 (SOI and temperature) and the respiratory release phase in vegetation (NDVI close to zero). The models include general control variables of a constant, a linear time trend and 11 centered seasonal dummy variables (annual frequency -1 dummies with mean zero in the long run), which are subject to selection so not automatically included in terminal models.

To capture interseasonal transport dynamics, interaction terms for winter/summer are included with each vegetation measure. Summer is defined as May–October and winter is defined as November–April. Thus binary variables (winter and summer weights) are interacted with all region-specific NDVI variables and included in the GUM. This could be extended to smoothed weights following a sine/cosine pattern, but binary weights are expected to capture the main seasonal effect of different atmospheric transport.

Sample 1 includes all variables measured at a monthly frequency: natural control variables as well as the first three components of industrial production. Sample 2 includes all short-term dynamics as in sample 1, and adds the first three low-frequency interpolated fossil fuel emission components.

Since ΔCO_2 is a highly autocorrelated series when measured monthly, a long lag length is allowed. Based on the partial-autocorrelation function (PACF) of ΔCO_2 , the longest lag length is selected to be 12, as longer lags fall below the 95 percent critical level. Figure 12.14 shows the PACF for ΔCO_2 . Additionally, 12 lags of each independent variable are added to GUM 1. Lag lengths for GUM 2 are specified based on selection from GUM 1 to compensate for a larger number of independent variables (see Section 6).

As controls for the terrestrial biosphere are included in every GUM, it is the NDVI measure that defines the maximum number of available observations. CO_2 is measured at Mauna Loa from 1958:3 until the present, but NDVI data are only available from 1981:7 until 2002:12, so the vegetation control limits the maximum number of observations to 258. A maximum lag length of 12 then leads to the estimated sample size being 246 observations (T = 246).

Impulse indicator saturation includes one binary variable for each observation. In the case of a model with 12 lags, this means that 246 individual binary variables are added

Table 12.6 GUM: sample 1

Variables included	Lag length
Temperature	12
NDVI PC1 Eurasia (Eur) + winter interaction	12
NDVI PC2 Eurasia (Eur) + winter interaction	12
NDVI PC3 Eurasia (Eur) + winter interaction	12
NDVI PC1 North America (NA) + summer interaction	12
NDVI PC2 North America (NA) + summer interaction	12
NDVI PC3 North America (NA) + summer interaction	12
SOI	12
Industrial production comp. 1	12
Industrial production comp. 2	12
Industrial production comp. 3	12
Constant	yes
Trend	yes
Centered seasonal variables	yes
Impulse indicators	yes
Total variables	492

to the GUM. Tables 12.6 and 12.7 provide an overview of the variables making up the sample-specific GUMs.

Non-stationarity

Non-stationarity is a common feature of time-series data. Broadly, a non-stationary process is defined as a process whose distribution changes over time, for example the mean or variance of the process is non-constant (see Hendry and Juselius, 2000, for a detailed discussion). The trending level of CO_2 measured at Mauna Loa is non-stationary as its mean is increasing with time. There are various forms of non-stationarity. The time series could be integrated processes (a time series is said to be integrated of order r, or I(r), if differencing the series r times yields a stationary process), or alternatively there could be structural breaks/shifts in coefficients or levels. Non-stationarity through structural breaks need not be removed by differencing. IIS is used to detect structural breaks (as well as mis-specification): if a large number of impulse indicators is selected, the model may be mis-specified, data badly mis-measured; or there are breaks in the data.

Non-stationary data are not a problem in automatic model selection so long as it is approached correctly. Selection in *Autometrics* is primarily based on *t*- and *F*-statistics that under non-stationarity can follow non-standard distributions. Sims et al. (1990) show that the limit distributions of these test statistics are standard if the original equation can be rewritten in terms of coefficients on mean-zero stationary variables. That is to say, the actual re-parametrization is not required, merely the existence of a linear re-parametrization in mean-zero variables is sufficient for the test statistics in the original equation to have standard distributions. So long as the equation can be written that way, selection based on *t*- and *F*-statistics will be valid. However, there are potential problems

Variables included	Lag length
Temperature	6
NDVI PC1 Eurasia (Eur) + winter interaction	12
NDVI PC2 Eurasia (Eur) + winter interaction	12
NDVI PC3 Eurasia (Eur) + winter interaction	12
NDVI PC1 North America (NA) + summer interaction	12
NDVI PC2 North America (NA) + summer interaction	12
NDVI PC3 North America (NA) + summer interaction	12
SOI	6
Industrial production comp. 1	6
Industrial production comp. 2	6
Industrial production comp. 3	6
CO ₂ emissions comp. 1	6
CO_2 emissions comp. 2	6
CO_2 emissions comp. 3	6
Constant	yes
Trend	yes
Centered seasonal variables	yes
Impulse indicators	yes
Total variables	483

Table 12.7GUM: sample 2

during selection when a path is considered in which this re-parametrization is not possible. Such a difficulty is hard to avoid, so is handled here by using tight significance levels that allow for possible non-standard distributions. As Figure 12.5 showed, $\Delta_{12}CO_2$ manifests a strong stochastic trend so is clearly not stationary: in monthly data, that trend is 'swamped' by the seasonal variation. Due to this strong seasonality, augmented Dickey–Fuller (Dickey and Fuller, 1981) type tests would not successfully identify a unit root. Seasonal unit-root tests are somewhat unreliable as other determinants, such as breaks, are not included, but based on the apparent stochastic trend in $\Delta_{12}CO_2$, it is safe to say that ΔCO_2 is integrated of order one, I(1). Unit roots can also not be rejected for most of the independent time-series variables that are included as potential determinants. The order of integration of the dependent variables is then the same as the order of integration of the independent variables. Given that the model can be written as coefficients on stationary mean-zero variables (see Banerjee et al., 1993), we proceed by estimating the model in levels with tight significance levels to account for selection effects where an I(0) re-parametrization is not possible.

Algorithm Specification

The *Autometrics* algorithm is used to estimate and select within the GUMs described in Tables 12.6 and 12.7. The algorithm is calibrated to the following parameters. Selection is done at a 0.1 percent significance level, considerably tighter than the conventional 5 percent or 1 percent used in the literature. In models starting with *K* irrelevant variables,

this implies that on average 0.001K irrelevant variables are retained in the terminal model. The current model is reduced by removing the least significant variable until variables cannot be dropped at the 0.1 percent level. At the termination of each path, models are backtested to the initially specified GUM when feasible. Diagnostic tests, defined formally below (see Doornik and Hendry, 2009), are conducted at a 1 percent level, for normality, heteroskedasticity, coefficient constancy (set to a 70 percent sample split), residual autocorrelation and autoregressive conditional heteroskedasticity, both based on eight lags. The specified GUM includes more variables than observations with IIS, so diagnostic tests are only applied to terminal models, and, if satisfactory, conventional standard errors are used.

6 RESULTS

Overview

The specifications given in Tables 12.6 and 12.7 are selected from general unrestricted models using *Autometrics* by the algorithm outlined in Section 5. Sample 1 includes all natural controls and monthly anthropogenic components; sample 2 includes all natural controls as well as monthly and annual anthropogenic components. The crucial feature of *Autometrics* is determining the selection of variables, rather than their estimated coefficients, although bias corrections have been implemented. Below, we also note the relative importance of individual variables through decomposition of explained variance.

The selection algorithm results in 14 terminal models for sample 1, and 16 terminal models for sample 2. It may surprise that so many congruent, undominated different representations can be found at such a tight significance level as 0.1 percent, but this merely reflects the high collinearity both between the different series and over time as represented by their lagged values. Most of the terminal models are minor variations on others as the final GUMs had 21 and 30 variables respectively.⁴ The final models are selected from the set of terminal models by the smallest SIC value. Equations (12.9) and (12.10) show the selected final models for sample 1 and sample 2 respectively.

$$\Delta CO_{2,t} = \underbrace{0.24}_{(0.053)} \Delta CO_{2,t-1} - \underbrace{0.67}_{(0.049)} \Delta CO_{2,t-2} + \underbrace{0.20}_{(0.05)} \Delta CO_{2,t-3}$$
(12.9)
$$- \underbrace{0.32}_{(0.037)} \Delta CO_{2,t-4} + \underbrace{0.24}_{(0.057)} IP_{1_{t-1}} - \underbrace{0.30}_{(0.056)} IP_{1_{t-4}}$$

$$- \underbrace{0.20}_{(0.034)} IP_{2_{t-4}} + \underbrace{0.15}_{(0.035)} IP_{3_{t}} + \underbrace{0.003}_{(0.0005)} Temp_{t-4}$$

$$- \underbrace{0.006}_{(0.0013)} SOI_{t-5} - \underbrace{0.042}_{(0.007)} NDVI_{1,Eur_{t-1}} + \underbrace{00.019}_{(0.003)} NDVI_{1,Eur_{t-10}}$$

$$- \underbrace{0.020}_{(0.006)} \underbrace{\text{NDVI}_{1,\text{Eur}_{t_{-1}^{-2}}}}_{(0.007)} + \underbrace{0.026}_{(0.007)} \underbrace{\text{w_NDVI}_{3,\text{Eur}_{t-8}}}_{(0.007)}$$

$$\hat{\sigma} = 0.212 \ T = 246 \ n = 14 \ \text{SIC} = -0.006 \ F_{ar}(8, 224) = 1.55$$

 $\chi^2_{nd}(2) = 1.47 \ F_{reset}(2, 230) = 1.42 \ F_{arch}(8, 230) = 1.07 \ F_{het}(28, 217) = 1.42$

Let F_{name} denote an approximate Lagrange-multiplier *F*-test, then F_{ar} tests for autocorrelation of order *k* (Godfrey, 1978), F_{het} tests for heteroscedasticity (White, 1980), F_{arch} tests for *k*th-order autoregressive conditional heteroskedasticity (ARCH: Engle, 1982), F_{reset} tests for functional-form mis-specification (White, 1980), and $\chi^2_{nd}(2)$ tests for non-normality (Doornik and Hansen, 2008).

$$\Delta CO_{2,t} = -\frac{0.57}{(0.051)} \Delta CO_{2,t-2} - \frac{0.22}{(0.050)} \Delta CO_{2,t-4} + \frac{00.27}{(0.062)} IP_{1_{t-1}}$$
(12.10)

$$-\frac{0.34}{(0.063)} IP_{1_{t-4}} - \frac{0.28}{(0.039)} IP_{2_{t-4}} + \frac{0.16}{(0.035)} IP_{3_t} - \frac{0.74}{(0.21)} I_{1990(7)}$$

$$+\frac{0.004}{(0.006)} Temp_{t-4} - \frac{0.007}{(0.001)} SOI_{t-5} - \frac{0.047}{(0.008)} NDVI_{1,Eur_t}$$

$$-\frac{0.044}{(0.006)} NDVI_{1,Eur_{t-3}} + \frac{0.033}{(0.006)} NDVI_{1,Eur_{t-11}} - \frac{0.042}{(0.009)} NDVI_{1,Eur_{t-12}}$$

$$+\frac{0.030}{(0.008)} NDVI_{2,Eur_{t-11}} - \frac{0.029}{(0.007)} w_NDVI_{2,Eur_{t-7}} - \frac{0.022}{(0.006)} w_NDVI_{2,Eur_{t-8}}$$

$$+\frac{0.035}{(0.007)} w_NDVI_{3,Eur_{t-8}} - \frac{0.039}{(0.008)} s_NDVI_{1,NA_{t-8}} + \frac{0.034}{(0.007)} s_NDVI_{1,NA_{t-9}}$$

$$\hat{\sigma} = 0.199 T = 246 n = 19 SIC = -0.048 F_{ar}(8,219) = 1.62$$

$$\chi^2_{nd}(2) = 0.044 F_{reset}(2,225) = 0.042 F_{arch}(8,230) = 0.57 F_{het}(36,208) = 1.03$$

First: as is to be expected from the theory, controls for natural factors are selected in both final models. Temperature anomalies enter the model with a positive coefficient, likely capturing the effect of oceanic uptake such that CO_2 increases with higher temperatures. Vegetation controls through the principal components of NDVI are selected in both models, as is the control for Southern Oscillation. However, a key finding is that in both terminal models, natural controls are insufficient to account for the variation in the change of atmospheric CO_2 . Anthropogenic factors captured through components of industrial output indices are consistently selected in both models. Selection of these components is highly consistent over the two models as the selected production components are identical in models 1 and 2. Selection of these is robust to the addition of emissions components which are not selected in sample 2, suggesting that the high-frequency measures provide a better approximation for anthropogenic emissions measured at Mauna Loa.

Second: most selected variables enter the model in lagged form. Only the third principal component of production (in samples 1 and 2) and the first component of Eurasian NDVI (in sample 2) have an estimated immediate effect on the growth of CO_2 . Most anthropogenic emissions and vegetation growth 'lead' measured atmospheric CO_2 by suggested time periods of 1 to 12 months. Relative to the initial sizes of the GUMs, few variables are retained, yet relative to the tight significance levels, many more are retained than could be attributed to chance (less than 1 on average).

316 Handbook on energy and climate change

Third: the final models appear to be well specified. As a result of the algorithm, all models pass tests for normality, heteroskedasticity, residual autocorrelation and autoregressive conditional heteroskedasticity. The number of selected indicators from IIS is low. There is only one indicator selected for 1990:7. This suggests that the model is correctly specified and there appear to be no real structural breaks or shifts in the change of atmospheric CO₂. No deterministic variables are selected: no constant, time trend or centered seasonals appear in the final models. This suggests that changes in CO₂ are well approximated by the selected final variables covering anthropogenic and natural factors. If constants are included post-selection (both not statistically different from zero in models 1 and 2), R^2 can be used as a rough measure of goodness of fit. Both final models exhibit a high goodness of fit ($R^2 = 0.974$ and = 0.978 respectively). This is not a straight result of selection as *Autometrics* does not directly maximize the goodness of fit. Moreover, R^2 measures should not be attributed much weight when assessing models, as there are preferred likelihood-based measures that also account for the number of parameters included.

GUM: Sample 1

Sample 1 in equation (12.9) covers all variables measured at a monthly frequency. *Autometrics* in sample 1 with 246 observations estimated 345 models reducing the number of explanatory variables from an initial 492 (split into initial six blocks) down to 14 in the final model. The final model passes the stationarity test on the residual, unit roots ranging from lags 1 to 12 are rejected at the 1 percent level using an ADF test. There are no impulse indicator variables selected in the final GUM. Together this provides evidence for a well-specified model that forms a stationary relationship.

Selection: neither the constant, the linear time trend nor centered seasonal variables feature in the final model, so that the seasonality and increase in the growth of CO_2 are explained by the anthropogenic and natural factors. All selected variables (apart from IP_{3_1}) enter the model as lags, suggesting a delay in the effect of CO_2 emissions/absorption and measurement at Mauna Loa. The longest lag on an anthropogenic component is four months.

The anthropogenic sources that are selected are all three principal components for industrial production at lag lengths ranging from immediate t to t - 4. Component 1 is selected at t - 1 and t - 4 with opposite signs, suggesting that it enters the model mainly as a difference. As these variables are principal components of production indices, the coefficients are not straightforward to interpret. The key result is the consistent selection, relative importance and lag length of these, rather than the signs of individual estimated coefficients.

The non-seasonally weighted principal components for NDVI are mostly selected, but only the Eurasian region is included. Given that the growth cycle is relatively similar in North America and Eurasia, this should not be over-interpreted. It is likely that the PCs for North American and Eurasian vegetation capture a very similar pattern and are to a considerable extent interchangeable. The negative coefficient on the first PC of vegetation (at t - 1, as well as in the long-run solution below) supports the theory that increased vegetation activity slows down CO₂ growth. The near equal magnitude, opposite signs on t - 10 and t - 12 suggest a difference, a pattern also seen in (12.10). The coefficient

Variable	Partial R^2	Hierarchical partitioning
$\Delta CO_{2,t-2}$	0.4472	0.0661
ΔCO_{2t-4}	0.2461	0.0643
NDVI _{1 Eur.}	0.1583	0.0257
IP_{2}	0.1259	0.0036
NDVI _{1 Fur}	0.1259	0.2795
$\operatorname{Temp}_{t-4}$	0.1225	0.0067
IP ₁	0.1105	0.0117
SOI_{t-5}	0.0836	0.0014
ΔCO_{2t-1}	0.0796	0.1954
IP ₃	0.0756	0.0074
$IP_{1,1}^{\gamma}$	0.0727	0.0157
$\Delta \dot{CO}_{2t-3}$	0.0619	0.0372
w_NDVI _{3 Eur}	0.0578	0.0200
NDVI _{1,Eur_{t-12}}	0.0483	0.2648

Table 12.8 GUM: sample 1 relative importance

on temperature is positive at a lag of (t - 4), likely capturing reduced oceanic absorption under higher temperature. Southern Oscillation enters the model at a lag of five months with a negative coefficient. Thus, during El Niño episodes (SOI < 0), growth in CO₂ appears to increase, in line with findings of other papers (see Bacastow, 1976; Keeling and Revelle, 1985).

To quantify and assess the relative importance of each regressor, we decompose the total explained variation in ΔCO_2 . Decomposition is not straightforward when independent variables are correlated. We use two measures, partial R^2 and hierarchical partitioning. The partial R^2 provides a measure of the marginal contribution to explained variation for a given variable, while holding other factors constant. In hierarchical partition of each variable over all potential orderings of the variables (see Kruskal, 1987, and method LMG in Groemping, 2007).⁵ Individual variance contributions sum to unity and can be interpreted as percentages. This yields values for individual variables that sum to the total explained variance (R^2). Table 12.8 ranks the selected variables by the partial R^2 and also reports relative importance based on hierarchical partitioning.

Based on this decomposition, the single largest (non-autoregressive) contribution comes from the Eurasian NDVI principal component of vegetation followed by industrial production. The anthropogenic components cumulatively explain a large fraction of the variation in atmospheric CO_2 . Perhaps surprisingly, both the SOI and temperature are ranked low based on hierarchical partitioning.

GUM: Sample 2

Sample 2 in (12.10) covers all variables measured at a monthly frequency as well as interpolated annual components for long-term anthropogenic emissions. Lag selection for sample 2 is based on selection in sample 1. Selection in model 1 results in a maximum lag of 4 on anthropogenic components and 5 for temperature and SOI. The longest lag

for vegetation was selected at t - 12. Therefore GUM 2 starts with a maximum lag of 6 for anthropogenic components, and a maximum lag of 12 for vegetation. *Autometrics* in sample 2 with 246 observations estimated 571 models, reducing the number of explanatory variables from an initial 483 (split into initial six blocks) down to 19 in the final model. Unit roots are rejected for the model residuals at the 1 percent level using ADF tests covering up to 12 lags. There is only a single impulse indicator selected (1990:7), suggesting no major breaks or mis-specification.

Selection: sample 2 results in a slightly higher number of selected variables in the final model relative to sample 1. In terms of robustness, selection is highly consistent relative to sample 1. Anthropogenic components, as well as temperature and Southern Oscillation, are selected identically to model 1. No deterministic terms are selected.

Anthropogenic emissions are captured solely by the principal components for industrial production: the low-frequency fossil fuel measures are not selected. These findings suggest that high-frequency industrial production provides a better measure for anthropogenic factors than interpolated fuel emissions. This may appear surprising given that anthropogenic emissions directly measure emitted CO_2 : however, this result likely stems from the annual frequency of CO_2 emissions that miss any seasonal component.

In terms of natural controls, the estimated coefficient on temperature is positive and that on Southern Oscillation is negative, both as in the previous model. The selection of vegetation variables moves towards seasonally weighted Eurasian indicators but also adds North American NDVI measures. The selection of seasonally weighted vegetation measures supports the theory of atmospheric transport. The shift in selection of vegetation variables is likely due to the collinearity in the NDVI measure resulting from a similar growth cycle in North America and Eurasia. The selected indicator for 1990:7 (July) suggests there was a reduction in growth relative to other years. Although the cause for this is not apparent in the data, it should be noted that such an indicator creates a step shift in the level of CO_2 , so may correspond to China's resurgent growth thereafter.

Table 12.9 ranks the selected variables by relative importance based on the partial R^2 . Anthropogenic components rank similarly to model 1, with IP_{2*t*-4} being the second largest (non-autoregressive) contributor to explained variation after Eurasian vegetation. However, the measure of relative importance matters – based on hierarchical partitioning, anthropogenic measures are ranked lower than vegetation, while still above temperature and SOI.

Comparisons between Sample 1 and Sample 2

GUM sample 1 and GUM sample 2 are estimated as robustness checks. Comparing *Autometrics* selection in samples 1 and 2, there is highly consistent selection of key variables. Industrial production, temperature, vegetation and Southern Oscillation are selected identically in both models. Anthropogenic emissions are consistently selected in the form of high-frequency industrial production rather than low-frequency annual emissions. Conversely, selection of regional NDVI varies slightly between models, which could be the result of a similar growing season and pattern in North America and Eurasia as measured through NDVI data. Differences in selection from sample 1 to sample 2 may seem surprising given that all vegetation variables in sample 1 are identi-

Variable	Partial R ²	Hierarchical partitioning
$\overline{\Delta CO_{2,t-2}}$	0.3582	0.0298
NDVI _{1.Eur.}	0.2147	0.0290
IP_{2t-4}	0.1814	0.0040
NDVI _{1.Eur.}	0.1376	0.1750
Temp _{t-4}	0.1349	0.0055
SOI _{t-5}	0.1149	0.0013
$IP_{1_{t-4}}$	0.1109	0.0067
NDVI _{1.Eur.}	0.1098	0.0950
w_NDVI _{3.Eur.}	0.0974	0.0122
s_NDVI _{1.NA}	0.0936	0.0550
s_NDVI	0.0869	0.0464
IP_{3t}	0.0840	0.0034
NDVI _{1.Eur}	0.0813	0.1740
$\Delta CO_{2,t-4}$	0.0805	0.0778
IP ₁₄₁	0.0792	0.0086
w_NDVI _{2.Eur}	0.0770	0.0780
w_NDVI _{2.Eur}	0.0605	0.1125
NDVI _{2.Eur.}	0.0555	0.0810
I ₁₉₉₀₍₇₎	0.0507	0.0018

Table 12.9 GUM: sample 2 relative importance

cal to those in sample 2. However, the PCs are only orthogonal within regions, and are highly correlated between regions, as Figure 12.7 shows. The result of different selection may be due to computational short cuts, namely dropping branches in the tree search after a model failed backtesting, and block partitioning in *Autometrics*, which adversely affect consistent selection.

Overall, based on *Autometrics* selection, natural factors such as vegetation, temperature and Southern Oscillation are necessary but not sufficient to explain changes in atmospheric CO₂ measured at Mauna Loa. Industrial production variables are consistently selected. Most estimated effects from selected variables affect ΔCO_2 with a lag, and there seem to be few or no structural breaks in the relationships being modeled. As a further robustness check, future work will involve applying the estimation method to other measurement stations such as Barrow, Alaska (Keeling et al., 2008). Atmospheric CO₂ at Barrow has been measured from 1974:2 until 2007:12 and displays higher amplitude and higher autocorrelation due to its location relative to Mauna Loa.

Using relatively few assumptions, automatic model selection with *Gets* modeling can provide a tool to successfully model complex relationships. Starting from a broad GUM that nests the LDGP theoretically (accounting for natural and anthropogenic sinks and sources), the analysis proceeds with an agnostic approach to determine the key factors in changes in atmospheric CO_2 . First, automatic model selection reduces the GUM to terminal models. Shortcomings of automatic model selection are computational issues in *Autometrics* selection because of short cuts and block partitioning. However, the terminal models appear congruent (acceptable diagnostic tests, few to no indicators or seasonal dummies selected) and are also supported by theoretical conclusions from the



Figure 12.15 Graphs of fitted and actual values in differences and levels

broader literature – natural effects are selected with the expected signs on coefficients. Key in the results is that additionally to the natural determinants selected, all terminal models include a large number of anthropogenic factors. The explained variance is decomposed to establish the relative importance of the variables. Using hierarchical partitioning and ranking based on partial correlations, the anthropogenic contribution to explained variation is quantifiable and found to be high.

From Changes to Levels

The models estimated as GUM sample 1 and GUM sample 2 describe the change in atmospheric CO₂, Δ CO₂. While most of the analysis has been focused on the change in CO₂, it is possible to recover the estimated level from the models 1 and 2. Level estimates for model 1 and model 2, based on CO_{2,t} = Δ CO_{2,t} + CO_{2,t-1} are given in Figure 12.15 and show the close fit. The next step is to attribute the components of the long-run explanation. To do so, we derive the relation after all dynamics from lagged variables have been solved out (the 'long-run solution': see Hendry, 1995). In a simple ADL(1) of the form

$$y_t = \lambda_1 x_{1,t} + \lambda_2 y_{t-1} + e_t \tag{12.11}$$

where $|\lambda_2| < 1$, the long-run conditional expected value is

$$E[y|x_1] = \lambda_1 x_1 / (1 - \lambda_2) = \beta_1 x_1$$
(12.12)

Then, based on the theoretical specifications given in equations (12.5)–(12.7), the solved estimated model is in the approximate form given in equation (12.13):

$$\Delta \text{CO}_{2,t} = \beta_1 x_{1,t} + \beta_2 x_{2,t} + \dots + \beta_q x_{q,t} + \varepsilon_t$$
(12.13)

This can be rewritten as

$$CO_{2,t} = CO_{2,t-1} + \beta_1 x_{1,t} + \beta_2 x_{2,t} + \dots + \beta_q x_{q,t} + \varepsilon_t$$
 (12.14)

Recursive substitution for $CO_{2,t-1}, CO_{2,t-2}, \ldots$ in equation (12.14) yields

$$CO_{2,t} = CO_{2,0} + \beta_1 \sum_{j=1}^{t} x_{1,j} + \beta_2 \sum_{j=1}^{t} x_{2,j} + \dots + \beta_q \sum_{j=1}^{t} x_{q,j} + \overline{e}_t$$
(12.15)

where $\overline{e}_t = \sum_{j=1}^t \varepsilon_j$

We divide the variables into two different groups: let *s* be the number of variables $x_{i,t}$ that have a stationary cumulative sum (non-trending) $\sum_{j=1}^{t} x_{i,j} \sim I(0)$, so that q - s variables have a non-stationary cumulative sum (trending) $\sum_{j=1}^{t} x_{i,j} \sim I(r)$, where r > 0. Equation (12.15) can then be expressed as

$$\operatorname{CO}_{2,t} = \operatorname{CO}_{2,0} + \mathbf{x}_{s,t}' \boldsymbol{\beta}_s + \mathbf{x}_{q-s,t}' \boldsymbol{\beta}_{q-s} + \bar{\boldsymbol{e}}_t$$
(12.16)

where $\mathbf{x}'_{s,t}$ and $\mathbf{x}'_{q-s,t}$ are $s \times 1$ and $(q-s) \times 1$ column vectors respectively with $\sum_{j=1}^{t} x_{i,j}$ as their row elements. Equation (12.16) implies that the trending level of CO₂ is a function of the cumulative sums of the stationary ($\mathbf{x}'_{s,t}$) and non-stationary ($\mathbf{x}'_{q-s,t}$) variables in our model. Stationary variables in $\mathbf{x}'_{s,t}$ by nature cannot drive the trend. Only explanatory variables with non-stationary cumulative sums, $\mathbf{x}'_{q-s,t}$, determine the trend. This provides a straightforward method of evaluating the underlying factors of the long-run trend.

Out of the selected variables in models 1 and 2, only a subset exhibits trending cumulative sums, which are all the anthropogenic factors and the temperature anomaly. Both natural controls of NDVI and SOI remain approximately stationary around zero over time. Importantly, neither final model includes a deterministic intercept or trend, which on summation would become a linear or a quadratic time trend. However, summed variables do not have a straightforward interpretation in the case of PCs of industrial output and temperature. The trending temperature anomaly is likely a mutually supporting feedback effect, as mentioned in Section 4. Overall, the trend in the levels of CO_2 is derived from the trends in the independent variables in both estimated models, so is driven primarily by the PCs and partly by temperature. Specifically, the empirical equivalents of (12.16) for model 1 are

$$\mathbf{x}_{s,t}' \boldsymbol{\beta}_{s} = -0.0038 \sum_{j=1}^{t} \text{SOI}_{j} + \left(-0.0275 \sum_{j=1}^{t} \text{NDVI}_{1,\text{Eur}_{j}} + 0.0169 \sum_{j=1}^{t} w_{\text{NDVI}_{3,\text{Eur}_{j}}} \right) (12.17)$$

$$\mathbf{x}_{q-s,t}' \boldsymbol{\beta}_{q-s} = -0.037 \sum_{j=1}^{t} \mathrm{IP}_{1_{j}} - 0.127 \sum_{j=1}^{t} \mathrm{IP}_{2_{j}} + 0.097 \sum_{j=1}^{t} \mathrm{IP}_{3_{j}} + 0.0018 \sum_{j=1}^{t} \mathrm{Temp}_{j} \quad (12.18)$$



Figure 12.16 Level of CO, and cumulative sums of anthropogenic and natural factors

Figure 12.16 shows the resulting coefficient-weighted cumulative sums of vegetation (NDVI_{*EUR*} principal components), SOI, and the combined anthropogenic components (IP₁ to IP₃ and temperature trend) for both models 1 and 2, together with the recorded level of CO₂. The industrial production components and temperature approximate the level of CO₂ well, marking a slight slowdown in the trend around 1991–93. Both cumulative stationary components vary over a small range, so contribute little to the long-run changes. Even though the model is estimated in net inflows to atmospheric CO₂, in re-parametrized form it can explain the long-run trend – and attributes it primarily to anthropogenic emissions. This is an outcome of the data analysis and is not enforced.

7 CONCLUSIONS

We identified anthropogenic contributions to atmospheric CO_2 measured at Mauna Loa using an automatic model selection algorithm. Traditionally, estimation of anthropogenic effects on CO_2 relied on *a priori* selection of variables, which may not be appropriate in complex relations, low-frequency measures of anthropogenic emissions, and decompositions of time series. Using *Autometrics* in a general to specific modeling approach allows for model selection with more variables than observations, stringent mis-specification testing and a more agnostic way of modeling complicated interactions. The algorithm is applied to model changes in atmospheric CO_2 controlling for natural as well as anthropogenic sinks and sources without *a priori* restrictions of the determinants. While they are not completely robust to initial specification, we find that natural factors such as vegetation, temperature and the Southern Oscillation are necessary, but not sufficient in explaining variation of atmospheric CO_2 . Industrial production components measured monthly are highly significant and consistently selected in the estimated models. A higher proportion of variation should be attributed to anthropogenic sources than has been previously the case in the literature. Producing congruent models, our methodology introduces *Gets* modeling through *Autometrics* as a useful tool in modeling complicated climate relationships.

NOTES

- * This research was supported in part by grants from the Open Society Foundation and the Oxford Martin School. We thank Roger Fouquet for useful comments and Qin Duo for helpful notes on Chinese production.
- 1. 1 ppm by volume of CO₂ in the atmosphere is approximately equal to 2.13 gigatons of carbon: see Clark (1982) and http://cdiac.ornl.gov/pns/convert.html.
- Chinese gross industrial output is only available from 1979 to 1999 with a change in measure thereafter. Inclusion of this series did not improve selection at the cost of a reduction in observations.
- See Marland et al. (2011) for a detailed listing of which countries are included. Eastern Europe includes Russia, Central Asia includes China, and India is covered by the Far East category.
- 4. Detailed results are available from the authors on request.
- 5. An intercept is forced to be retained in selection of these models, though it is not statistically different from zero for models 1 and 2.

REFERENCES

- Australian Bureau of Meteorology (2011), 'Southern Oscillation index archive', http://www.bom.gov.au/ climate/current/soihtm1.shtml.
- Bacastow, R.B. (1976), 'Modulation of atmospheric carbon dioxide by the Southern Oscillation', *Nature*, **261**, 116–18.
- Bacastow, R.B., Keeling, C.D. and Whorf, T.P. (1985), 'Seasonal amplitude increase in atmospheric CO₂ concentration at Mauna Loa, Hawaii 1959–1982', *Journal of Geophysical Research*, 90, 10529–40.
- Banerjee, A., Dolado, J.J., Galbraith, J.W. and Hendry, D.F. (1993), Co-integration, Error Correction and the Econometric Analysis of Non-Stationary Data, Oxford: Oxford University Press.
- Bårdsen, G., Hurn, S. and McHugh, Z. (2010), 'Asymmetric unemployment rate dynamics in Australia', Creates Research Paper 2010-2, Aarhus University, Denmark.
- Bontemps, C. and Mizon, G.E. (2008), 'Encompassing: concepts and implementation', Oxford Bulletin of Economics and Statistics, 70, 721–50.
- Buermann, W., Lintner, B.R., Koven, C.D., Angert, A., Pinzon, J.E., Tucker, C.J. and Fung, I.Y. (2007), 'The changing carbon cycle at Mauna Loa observatory', *Proceedings of the National Academy of Sciences of the United States of America*, 104, 4249–54.
- Campos, J., Ericsson, N.R. and Hendry, D.F. (2005), Editors' introduction, in J. Campos, N.R. Ericsson and D.F. Hendry (eds), *Readings on General-to-Specific Modeling*, Cheltenham, UK and Northampton, MA, USA: Edward Elgar, pp. 1–18.
- Castle, J.L., Doornik, J.A. and Hendry, D.F. (2011), 'Evaluating automatic model selection', *Journal of Time Series Econometrics*, 3(1), DOI: 10.2202/1941–1928.1097.
- Castle, J.L., Doornik, J.A. and Hendry, D.F. (2012a), 'Model selection in equations with many "small" effects', *Oxford Bulletin of Economics and Statistics*, doi:10.1111/j.1468-0084.2012.00727.x.
- Castle, J.L., Doornik, J.A. and Hendry, D.F. (2012b), 'Model selection when there are multiple breaks', *Journal of Econometrics*, **169**, 239-46.
- Castle, J.L. and Hendry, D.F. (2011), 'Automatic selection of non-linear models', in L. Wang, H. Garnier and T. Jackman (eds), *System Identification, Environmental Modelling and Control*, New York: Springer, pp.229–50.
- Castle, J.L. and Shephard, N. (eds) (2009), *The Methodology and Practice of Econometrics*, Oxford: Oxford University Press.
- Chow, G.C. (1960), 'Tests of equality between sets of coefficients in two linear regressions', *Econometrica*, **28**, 591–605.

- Christopher, S.L., Feely, R.A., Gruber, N., Key, R.M., Lee, K., Bullister, J.L., Wanninkhof, R., Wong, C.S., Wallace, D.W.R., Tilbrook, B., Millero, F.J., Peng, T.H., Kozyr, A., Ono, T. and Rios, A.F. (2007), 'The oceanic sink for anthropogenic CO₂', *Science*, **305**, 367–71.
- Clark, W.C. (ed.)(1982), Carbon Dioxide Review, New York: Oxford University Press.
- Dickey, D.A. and Fuller, W.A. (1981), 'Likelihood ratio statistics for autoregressive time series with a unit root', *Econometrica*, **49**, 1057–72.
- Doornik, J.A. (2008), 'Encompassing and automatic model selection', Oxford Bulletin of Economics and Statistics, 70, 915–25.
- Doornik, J.A. (2009), 'Autometrics', in Castle and Shephard (2009), pp. 88-121.
- Doornik, J.A. (2010), 'Econometric model selection with more variables than observations', Working Paper, Economics Department, University of Oxford.
- Doornik, J.A. and Hansen, H. (2008), 'An omnibus test for univariate and multivariate normality', Oxford Bulletin of Economics and Statistics, 70, 927–39.
- Doornik, J.A. and Hendry, D.F. (2009), *Empirical Econometric Modelling using PcGive: Volume I*, London: Timberlake Consultants Press.
- Engle, R.F. (1982), 'Autoregressive conditional heteroscedasticity, with estimates of the variance of United Kingdom inflation', *Econometrica*, **50**, 987–1007.
- Enting, I.G. (1987), 'The interannual variation in the seasonal cycle of carbon dioxide concentration at Mauna Loa', *Journal of Geophysical Research*, 92, 5497–504.
- Erickson, D.J., Mills, R.T., Gregg, J., Blasing, T.J., Hoffmann, F.M., Andres, R.J., Devries, M., Zhu and Kawa, S.R. (2008), 'An estimate of monthly global emissions of anthropogenic CO₂: impact on the seasonal cycle of atmospheric CO₂', *Journal of Geophysical Research*, **113**.
- Federal Reserve (2011), 'Industrial production and capacity utilization', http://www.federalreserve.gov/releases/g17.
- Francey, R.J., Tanis, P.P., Allison, C.E., Enting, I.G., White, J.W.C. and Trolier, M. (1995), 'Changes in oceanic and terrestrial carbon uptake since 1982', *Nature*, 373, 326–30.
- Gerlach, T.M. (2011), 'Changes in oceanic and terrestrial carbon uptake since 1982', EOS, Transactions, American Geophysical Union, 92, 201–2.
- Godfrey, L.G. (1978), 'Testing for higher order serial correlation in regression equations when the regressors include lagged dependent variables', *Econometrica*, **46**, 1303–13.
- Granados, J.A.T., Ionides, E.L. and Carpintero, O. (2009), 'A threatening link between world economic growth and atmospheric CO, concentrations', Working Paper, University of Michigan.
- Granados, J.A.T., Ionides, E.L. and Carpintero, O. (2012), 'Climate change and the world economy: short-run and long-run determinants of amtospheric CO₂', *Environmental Science and Policy*, **21**, 50–62.
- Groemping, U. (2007), 'Estimators of relative importance in linear regression based on variance decomposition', American Statistician, 61(2), 139–47.
- Hansen, J. and Lebedeff, S. (1987), 'Global trends of measured surface air temperature', *Journal of Geophysical Research*, 92(13), 345–72.
- Hansen, J., Ruedy, R., Sato, M. and Lo, K. (2010), 'Global surface temperature change', *Review of Geophysics*, **48**, 2010RG000345.
- Hansen, J., Russel, G., Rind, D., Stone, P., Lacis, A., Lebedeff, S., Ruedy, R. and Travis, L. (1983), 'Efficient three-dimensional global models for climate studies: Models I and II', *Monthly Weather Review*, 111, 609–62.
- Hansen, J., Sato, M., Kharecha, P., Beerling, D., Berner, R., Masson-Delmotte, V, Pagani, M., Raymo, M., Royer, D.L. and Zachos, J.C. (2008), 'Target atmospheric CO₂: where should humanity aim?', *The Open Atmospheric Science Journal*, 2, 217–31.
- Hendry, D.F. (1995), Dynamic Econometrics, Oxford: Oxford University Press.
- Hendry, D.F. (2009), 'The methodology of empirical econometric modeling: applied econometrics through the looking-glass', in T.C. Mills and K.D. Patterson (eds), *Palgrave Handbook of Econometrics*, Basingstoke: Palgrave Macmillan, pp. 3–67.
- Hendry, D.F. (2011), 'Climate change: possible lessons for our future from the distant past', in S. Dietz, J. Michie and C. Oughton (eds), *The Political Economy of the Environment*, London: Routledge, pp. 19–43.
- Hendry, D.F. and Johansen, S. (2013), 'Model discovery and Trygve Haavelmo's legacy', *Econometric Theory*, forthcoming.
- Hendry, D.F., Johansen, S. and Santos, C. (2008), 'Automatic selection of indicators in a fully saturated regression', *Computational Statistics & Data Analysis*, 33, 317–35.
- Hendry, D.F. and Juselius, K. (2000), 'Explaining cointegration analysis: Part I', Energy Journal, 21, 1-42.
- Hendry, D.F. and Krolzig, H.-M. (2003), 'New developments in automatic general-to-specific modelling', in B.P. Stigum (ed.), *Econometrics and the Philosophy of Economics*, Princeton, NJ: Princeton University Press, pp. 379–419.
- Hendry, D.F. and Krolzig, H.-M. (2005), 'The properties of automatic Gets modelling', *Economic Journal*, 115, C32–C61.

- Hendry, D.F. and Mizon, G.E. (2011), 'Econometric modelling of time series with outlying observations', Journal of Time Series Econometrics, 3(1), DOI: 10.2202/1941–1928.1100.
- Hendry, D.F. and Richard, J.-F. (1989), 'Recent developments in the theory of encompassing', in B. Cornet and H. Tulkens (eds), *Contributions to Operations Research and Economics. The XXth Anniversary of CORE*, Cambridge, MA: MIT Press, pp. 393–440.
- Hendry, D.F. and Santos, C. (2010), 'An automatic test of super exogeneity', in M.W. Watson, T. Bollerslev and J. Russell (eds), *Volatility and Time Series Econometrics*, Oxford: Oxford University Press, pp. 164–93.
- Hoffman, P.F. and Schrag, D.P. (2000), 'Snowball Earth', Scientific American, 282, 68-75.
- Hofman, D.J., Butler, J.H. and Tans, P.P. (2009), 'A new look at atmospheric carbon dioxide', Atmospheric Environment, 43, 2084–86.
- Johansen, S. and Nielsen, B. (2009), 'An analysis of the indicator saturation estimator as a robust regression estimator', in Castle and Shephard (2009), pp. 1–36.
- Jones, C.D. and Cox, P.M. (2005), 'On the significance of atmospheric CO₂ growth rate anomalies in 2002–2003', *Journal of Geophysical Research*, **32**.
- Juselius, K. and Kaufmann, R. (2009), 'Long-run relationships among temperature, CO₂, methane, ice and dust over the last 420,000 years: cointegration analysis of the Vostok ice core data', *IOP Conference Series: Earth and Environmental Science*, 6.
- Kaufmann, R.K., Paletta, L.F., Tian, H.Q., Myeni, R.B. and D'Arrigo, R.D. (2008), 'The power of monitoring stations and a CO₂ fertilization effect: evidence from causal relationships between NDVI and carbon dioxide', *Earth Interactions*, **12**, 1–23.
- Kawa, S.R., Erickson, D.J., Pawson, S. and Zhu, Z. (2004), 'Global CO₂ transport simulations using meteorological data from NASA data assimilation system', *Journal of Geophysical Research*, 109, D18312.
- Keeling, C.D. (1973), 'Industrial production of carbon dioxide from fossil fuels and limestone', *Tellus*, **25**(2), 174–98.
- Keeling, C.D., Bacastow, R.B., Brainbridge, A.E., Ekdahl, C.A., Guenther, J.P.R. and Waterman, L.S. (1976), 'Atmospheric carbon dioxide variations at Mauna Loa observatory, Hawaii', *Tellus*, **6**, 538–51.
- Keeling, C.D., Chin, J.F. and Whorf, T.P. (1996), 'Increased activity of northern vegetation inferred from atmospheric CO, measurements', *Nature*, 382, 146–49.
- Keeling, C.D. and Revelle, R. (1985), 'Effects of El Niño/southern oscillation on the atmospheric content of carbon dioxide', *Meteoritics*, 20(2), 437–50.
- Keeling, C.D., Whorf, T.P., Whalen, M. and van der Plicht, J. (1995), 'Interannual extremes in rate rise of atmospheric carbon dioxide since 1980', *Nature*, 375, 666–70.
- Keeling, R.F., Piper, S.C, Bollenbacher, A.F. and Walker, S.J. (2008), 'Atmospheric CO₂ records from sites in the SIO air sampling network', in *Trends: A Compendium of Data on Global Change*, Carbon Dioxide Information Analysis Center, US Department of Energy, Oak Ridge, TN, USA, http://cdiac.ornl.gov/ trends/co2/sio-bar.html.
- Kohlmaier, G.H., Sire, E.O., Janecek, A., Keeling, C.D., Piper, S.C, and Revelle, R.(1989), 'Modelling the seasonal contribution of a CO₂ fertilization effect of the terrestrial vegetation to the amplitude increase in atmospheric CO₂ at Mauna Loa observatory', *Tellus*, **41B**, 487–510.
- Kruskal, W. (1987), 'Relative importance by averaging over orderings', American Statistician, 41(1), 6-10.
- Levin, I., Graul, R. and Trivett, N.B.A. (1995), 'Long-term observations of atmospheric CO₂ and carbon isotopes at continental sites in Germany', *Tellus*, 47B, 23–34.
- Lucht, W., Prentice, I.C., Myeni, R.B., Sitch, S., Friedlingstein, P., Cramer, W., Bousquet, P., Buermann, W. and Smith, B. (2002), 'Climate control of the high-latitude vegetation greening trend and Pinatubo effect', *Science*, 296, 1687–9.
- Marland, G., Boden, T.A. and Andres, R. (2011), 'Global, regional, and national fossil fuel CO₂ emissions', in *Trends: A Compendium of Data on Global Change*, Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, US Department of Energy, Oak Ridge, TN, USA, http://cdiac.ornl.gov/trends/ emis/overview.html.
- Marland, G. and Rotty, R.M. (1984), 'Carbon dioxide emissions from fossil fuels: a procedure for estimation and results for 1950–1982', *Tellus*, 36B, 232–61.
- Ministry of Statistics Government of India (2011), 'Index of industrial production (IIP)', http://mospi.nic.in/.
- Mizon, G.E. and Richard, J.F. (1986), 'The encompassing principle and its application to non-nested hypothesis tests', *Econometrica*, 54, 657–78.
- Myeni, R.B., Hall, F.G., Sellers, P.J. and Marshak, A.L. (1995), 'The interpretation of spectral vegetation indexes', *Transactions on Geoscience and Remote Sensing*, 33(2), 481–6.
- NASA Goddard Institute for Space Studies (GISS) (2011), 'GISS surface temperature analysis', http://data.giss.nasa.gov/gistemp/.
- Nevison, C.D., Mahowald, N.M., Doney, S.C, Lima, I.D., van der Werf, G.R., Randerson, J.T., Baker, D.F., Kasibhatla, P. and McKinley, G.A. (2008), 'Contribution of ocean, fossil fuel, land biosphere, and biomass

burning carbon fluxes to seasonal and interannual variability in atmospheric CO₂', *Journal of Geophysical Research*, **113**.

- Newell, N.D. and Marcus, L. (1987), 'Carbon dioxide and people', *Society for Sedimentary Geology*, **2**(1), 101–3.
- OECD (2011), 'Composite leading indicators: MEI', http://stats.oecd.org/Index.aspx.
- Office of National Statistics (2011), 'Primary production', http://www.statistics.gov.uk/hub/business-energy/.
- Orr, C.J., Maier-Reimer, E., Mikolajewicz, U., Monfray, P., Sarmiento, J.L., Toggweiler, J.R., Taylor, N.K., Palmer, J., Gruber, N., Sabine, C.L., Le Quere, C., Key, R.M. and Boutin, J. (2001), 'Estimates of anthropogenic carbon uptake from four three-dimensional global ocean models', *Global Biogeochemical Cycles*, 15(1), 43–60.
- Randerson, T.J., Thompson, M.V., Conway, T.J., Fung, I.Y. and Field, C.B. (1997), 'The contribution of terrestrial sources and sinks to trends in the seasonal cycle of atmospheric carbon dioxide', *Global Biogeochemical Cycles*, 11(4), 535–60.
- Ritschard, R.L. (1992), 'Marine algae as a CO, sink', Water, Air and Soil Pollution, 64, 289-303.
- Ruddiman, W. (ed.)(2005), Plows, Plagues and Petroleum: How Humans took Control of Climate, Princeton, NJ: Princeton University Press.
- Salkever, D.S. (1976), 'The use of dummy variables to compute predictions, prediction errors and confidence intervals', *Journal of Econometrics*, 4, 393–7.
- Schwarz, G. (1978), 'Estimating the dimension of a model', Annals of Statistics, 6, 461-4.
- Sims, C.A., Stock, J.H. and Watson, M.W. (1990), 'Inference in linear time series models with some unit roots', Econometrica, 58, 113–44.
- Sundquist, E.T. and Keeling, R.F. (2009), 'The Mauna Loa carbon dioxide record: lessons for long-term earth observations', *Geophysical Monograph Series*, 183, 27–35.
- Taguchi, S., Murayama, S. and Higuchi, K. (2003), 'Sensitivity of inter-annual variation of CO₂ seasonal cycle at Mauna Loa to atmospheric transport', *Tellus*, 55B, 547–54.
- Tans, P. and Keeling, R. (2011), 'Mauna Loa, monthly mean carbon dioxide', Scripps Institution of Oceanography (scrippsco2.ucsd.edu/) and NOAA/ESRL http://www.esrl.noaa.gov/gmd/ccgg/trends/.
- Thoning, K.W. and Tans, P.P. (1989), 'Atmospheric carbon dioxide at Mauna Loa observatory 2. Analysis of the NOAA GMCC data, 1974–1985', Journal of Geophysical Research, 94, 8549–65.
- Troup, A.J. (1965), 'The "Southern Oscillation", Quarterly Journal of the Royal Meteorological Society, 91, 490–506.
- Tucker, C.J., Pinzon, J., Brown, M. and GIMMS/GSFC/NASA (2010), 'ISLSCP II GIMMS monthly NDVI, 1981–2002', in Forrest G. Hall, G. Collatz, B. Meeson, S. Los, E. Brown de Colstoun and D. Landis (eds), ISLSCP Initiative II Collection. Data set, http://daac.ornl.gov/ from Oak Ridge National Laboratory Distributed Active Archive Center, Oak Ridge, TN, USA.
- Watson, A.J., Nightingale, P.D. and Cooper, D.J. (1995), 'Modeling atmosphere ocean CO₂ transfer', *Philosophical Transactions of the Royal Society of London Series B – Biological Sciences*, 353, 41–51.
- White, H. (1980), 'A heteroskedastic-consistent covariance matrix estimator and a direct test for heteroskedasticity', *Econometrica*, 48, 817–38.
- Wigley, T.M.L. (1983), 'The pre-industrial carbon dioxide level', Climate Change, 5, 315-20.
- Worrell, E., Price, L., Martin, N., Hendriks, C. and Meida, L.O. (2001), 'Carbon dioxide emissions from the global cement industry', *Annual Review of Energy and the Environment*, 26, 303–29.