Atmospheric CO2 and global temperatures: the strength and nature of their dependence.

Granville Tunnicliffe Wilson*

May 3, 2015

Abstract

There is now considerable scientific consensus that the acknowledged increase in global temperatures is due to the increasing levels of atmospheric carbon dioxide arising from the burning of fossil fuels. Large scale global circulation models support this consensus and there have also been statistical studies which relate the trend in temperatures to the carbon dioxide increase. However, causal dependence of one trending series upon another cannot be readily proved using statistical means. In this paper we model the trend corrected series by times series methods which provide a plausible representation of their dependence. A consequence of trend correction and our use of relatively short series is that our model is unable to give precise long-term predictions, but it does illuminate the relationships and interaction between the series.

Keywords: Time series prediction, spectral coherency, structural VAR, graphical modeling

1 The series

In a previous unpublished conference paper we modeled three series: (i) the atmospheric carbon dioxide (CO2) concentration in parts per million (ppm) observed at Mauna Loa, (ii) the annual global mean temperature anomaly known as HadCRUT3 and (iii) the southern oscillation index (SOI). The HadCRUT3 series is a combination of the sea surface temperature (SST) anomaly series known as HadSST-gl and the land surface temperature (LST) anomaly series known as CRUTEM3. In the present paper we use these two separate series (with CRUTEM3 updated to CRUTEM4) in place of the combined series, which leads to simplification of the model.

The sources for the series are

(i) CO2: ftp://aftp.cmdl.noaa.gov/products/trends/co2/

(ii) HadSST-gl and CRUTEM4: http://www.cru.uea.ac.uk/cru/data/temperature/

(iii) SOI: http://www.cpc.ncep.noaa.gov/data/indices/

*Department of Mathematics and Statistics, Lancaster University, Lancaster LA1 4YF, UK.

The SOI is the observed sea level pressure difference between Tahiti and Darwin, Australia, and is strongly related to ocean temperatures across the Eastern Pacific Ocean. In our models we should consider the SOI series to be a proxy variable for these ocean temperatures, or possibly ocean temperature gradients across the region, because, unlike the temperature series, it has no visually evident trend. It is formed as the difference of the two standardized series, but to minimize preprocessing of the data we omit the standardization and we will also reverse the sign of the difference so our series is positively correlated with sea surface temperatures.

The CO2 and SOI series are precisely defined and measured at respectively one and two stations. In contrast, the land and sea surface temperature anomalies are compiled from a large number of measurements around the globe. We use the global values but they are also available separately for the northern and southern hemispheres. However, we will be modeling annual mean values from 1959, the first full year of Mauna Loa records, to 2014, giving series of length 56. With 6 series, saturated forms of the vector autoregressive (VAR) models that we use would have an excessive number of parameters for models of more than very low order. We shall show in the next section that the spectral coherency between all the series is high, suggesting that there is little need for more information.

We model the annual series shown in Figure 1 because of the substantial within-year variability of CO2 and temperatures due to the natural seasonal influences.

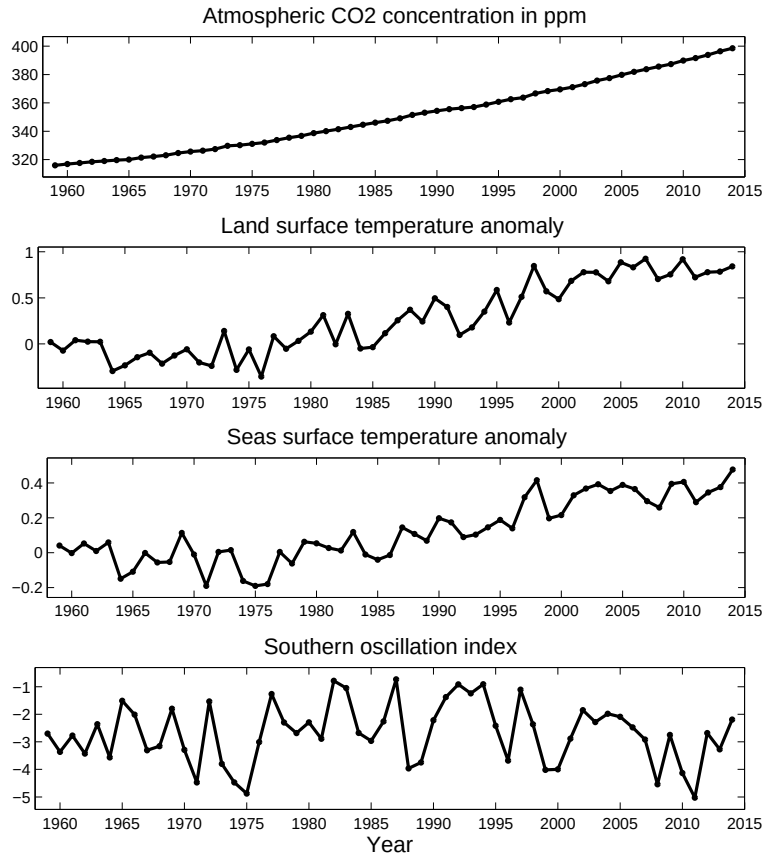

Figure 1: The four climate series analyzed and modeled in this paper.

Within one year we expect each variable to influence others, for example CO2 levels influence air temperature and uptake of CO2 in the oceans depends on sea surface temperatures. We will therefore develop structural forms of vector autoregressions with simultaneous equation relationships between the innovations in current annual values, to represent the net effect of these mutual influences.

2 Spectral coherency between the series

Trends are evident in the raw series apart from the SOI. A time series model for the trending carbon dioxide and global temperature series over a longer time period is given by Young (2014). However, our intention as expressed in the abstract is to correct for these trends for which we use ordinary least squares (OLS) regression. The CO2 series is corrected for a quadratic trend and the temperature series for linear trends which are visually evident. The SOI series is simply mean corrected; trend correction makes very little difference. Figure 2 shows these trend corrected series. The land and sea temperatures look very similar, but they have no obvious similarity with the other series.

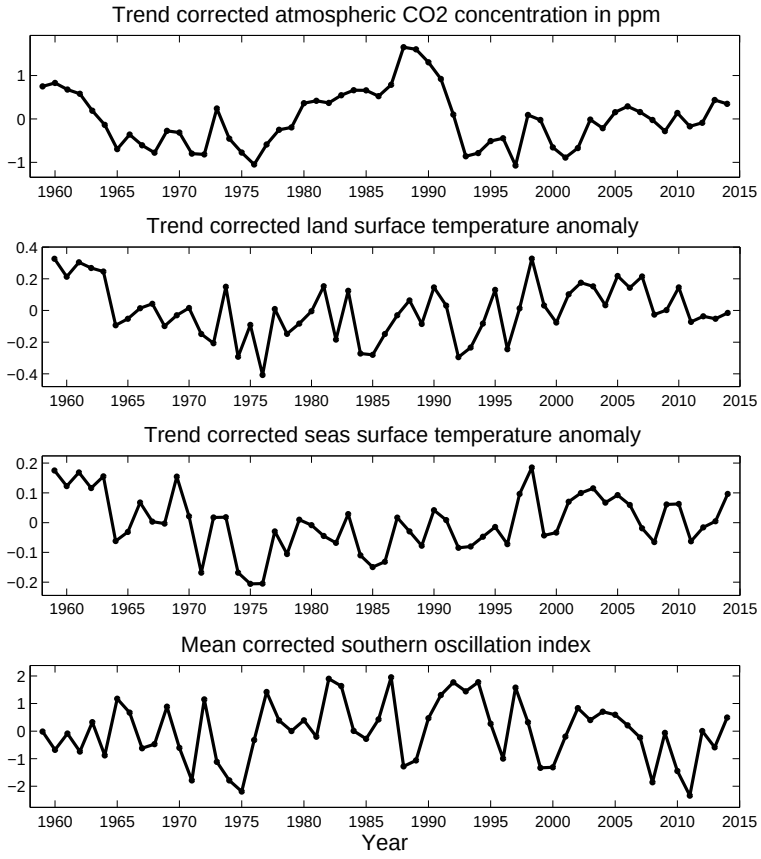

Figure 2: The four climate series after correction for mean and trend components.

However, the spectral coherency between all the series, which takes into account linear lagged dependence, is quite significant, as shown in the upper half of Figure 3, in a form

introduced by Dahlhaus (2000, p.167).

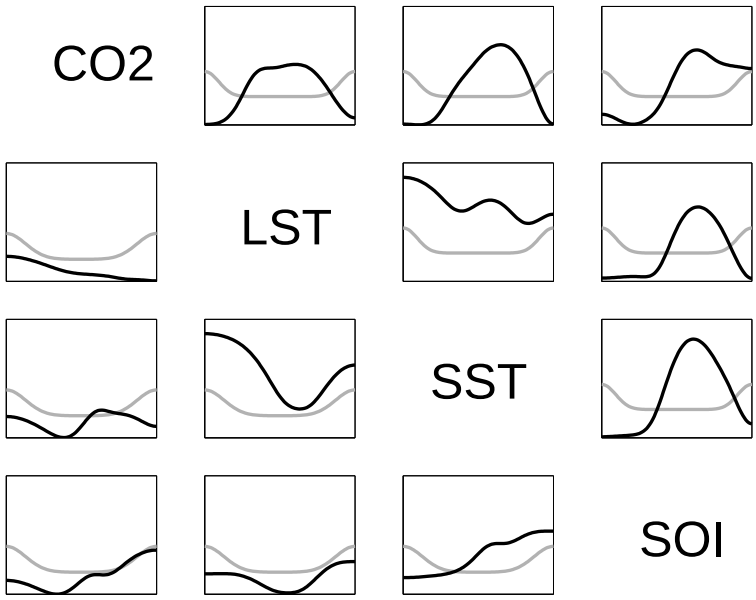

Figure 3: An array of plots showing the pairwise squared coherencies above the diagonal and partial coherencies below the diagonal, between the four climate series indicated on the diagonal. The squared coherencies are smoothed using a bandwidth of 0.2, and the partial coherencies using a bandwidth of 0.3. The coherencies are shown in the solid black line and their 95% significance limits by the solid gray line.

The lower half of Figure 3 shows the partial coherencies which measure, for a pair of series, the further dependence of the one upon the second, given its dependence upon the other two. These are seen to be much weaker than the pairwise coherencies. We summarize the partial coherencies in Figure 4 in which the links between the series reflect the significance observed in the lower half of Figure 3. A solid line corresponds to significant coherency over some part of the frequency range and a broken line corresponds to marginal significance, otherwise no link is shown.

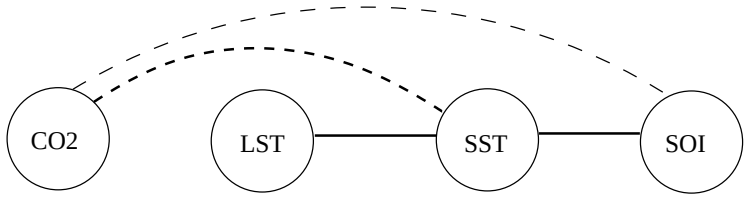

Figure 4: The partial coherency graph between the four climate series.

This graph provides only a limited description of the dependence between the series but it can be related to their causal dependence as described by Dahlhaus and Eichler (2003). Subject, of course, to statistical uncertainty, the graph implies that in a structural VAR model for the series there should be no explicit dependence of either CO2 or LST on present

or past values of the other. We will comment further on this after we have built such a model.

Spectral analysis can also estimate the lagged response of one series to another, and we show these for selected pairs of the climate series in Figure 5. Thus the uppermost plot of this figure shows the regression coefficients of SOI on SST at positive and negative lags. At positive lags the coefficients describe the effect of the current value of SST on future values of SOI - the *response* of SOI to SST. The regressions are estimated, however, in the presence of feedback between the series, so the coefficients do not correctly estimate the causal response, and significant coefficients may be observed at negative lags. For the correct causal response we will use a VAR model.

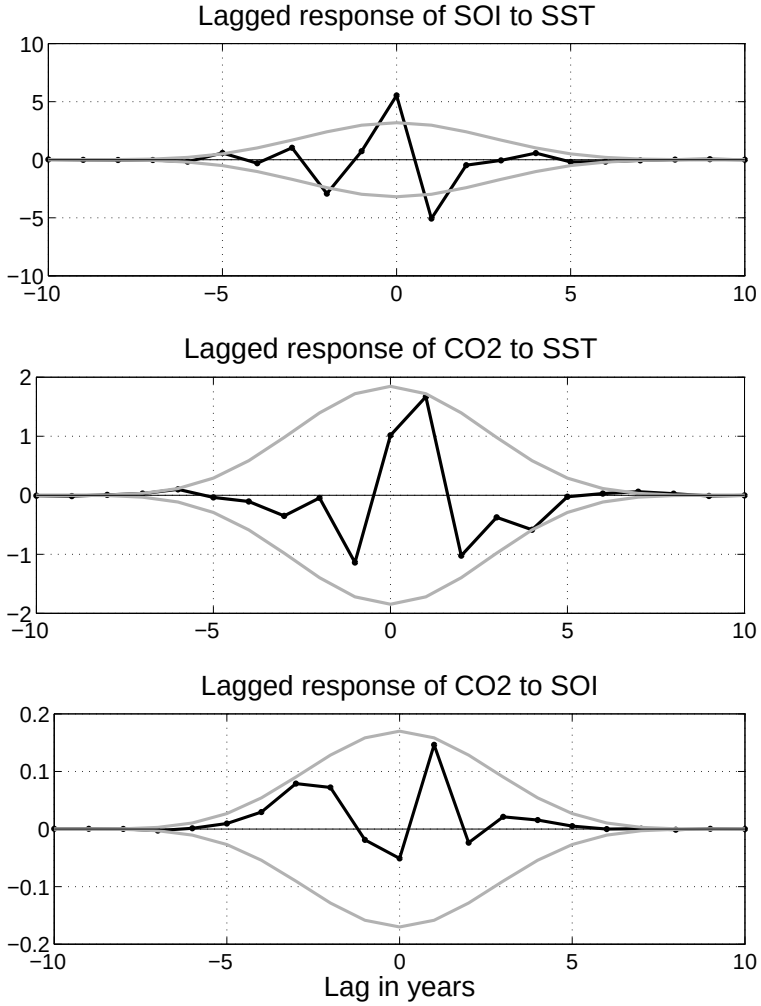

Figure 5: Lagged responses of selected pairs of climate series.

The response of SOI to SST shows the acknowledged strong positive relationship (with our sign convention) in the same year, but this is followed by a strong negative relationship in the subsequent year. This reflects the eponymous oscillatory nature of the SOI, with a period of between two and three years. The response of CO2 to SOI is positive in the same and subsequent year, though only marginally significant when estimated by spectral means.

This reflects the fact that at higher temperatures the sea cannot absorb so much CO2, so there is an apparent consequent net increase. Of particular interest is the response of CO2 o SOI at a lag of one year. A low SOI is associated with cooler sea temperatures in the Eastern Pacific with upwelling of cold seawater close to the South American west coast. A high SOI (with our sign convention) is associated with the El Niño effect in which this upwelling fails and warmer water floods into that area. This warmer water can absorb much less CO2 than than the cold upwelling. The response shown in Figure 5 suggests that actually our annual SOI series is predictive of this effect by one year.

Spectral estimation of these responses are of course limited in their causal interpretation and as semi-parametric estimators they are not as statistically efficient as the parametric models we consider next. They do, however, give useful predictive information and insight into the relationships between the series.

3 A standard VAR model for the series.

As a preliminary step in developing an empirically identified structural VAR model we first identify a saturated VAR model of order p, in the standard form which may be found, for example in Lütkepohl (1993) or Reinsel (1993):

$$x_t = \Phi_1 x_{t-1} + \Phi_2 x_{t-2} + \cdots + \Phi_p x_{t-p} + e_t, \tag{1}$$

where x_t is the vector of series values at time t and $\Phi_1, \ldots, \Phi_p$ are the matrix coefficients of dependence upon lagged series values. The error or innovation vector e_t has, in general, correlated elements with covariance matrix V_e. It is also multivariate white noise and uncorrelated with all past observations x_{t-k} for $k > 0$.

The series are all zero mean as a consequence of mean and trend correction and we assume they are second order stationary. The VAR model is therefore fitted for increasing orders of p and the AIC, Akaike (1973), plotted as in the left hand plot of Figure 6. The model is fitted by exact maximum likelihood under the normal assumption. This gives very similar estimates to the use of lagged OLS regression but uses information in the full length of series, which is preferable for short series. Lagged regression of order k necessarily uses the reduced series length of $n - k$. The plot on the right of Figure 6 shows the AIC as modified by Hurvich and Tsai (1989) to improve the order selection by this criterion in small samples. In fact both have their minimum for the model of order 2, though the modified AIC is only marginally less than its value at order 1. For univariate series the AIC is known not to be consistent for order selection, with a small probability of selection of an order somewhat higher than the true one. However, the situation is not so simple for our multivariate series because 16 new coefficients are introduced when the order is increased by 1. The probability of overestimating the true order, p, is then quite low, and on the contrary, in small samples the true order may be underestimated if only a small number of the 16 coefficients are non-zero at lag p. In fact, for the model of order 2 the dependence of LST on SOI at lag 2 has t-value 2.00 and the dependence of SST on LST at lag 2 has t-value 2.30. This supports the selection of the order $p = 2$ indicated by the modified AIC.

We now demonstrate the extent to which this VAR(2) model for the four climate series captures their dependence by using it to estimate the whole record of the SST series from

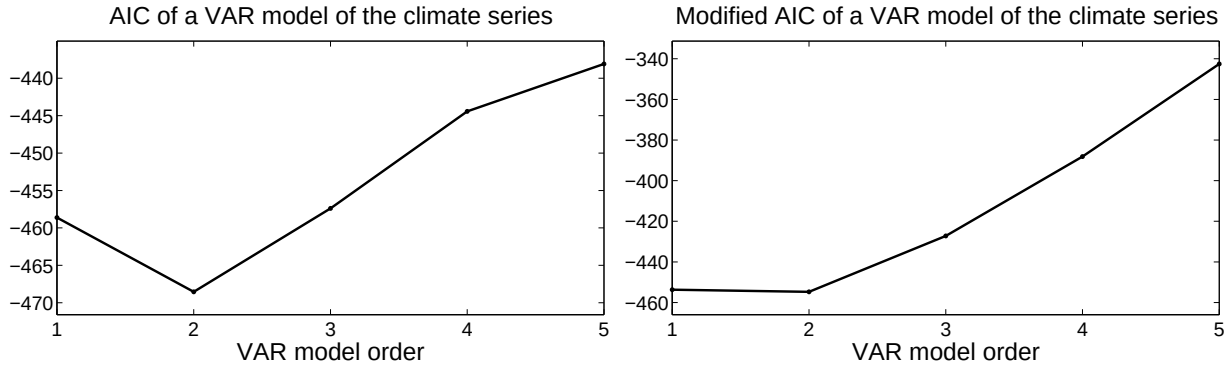

Figure 6: Plots of the AIC and modified AIC of a VAR model for the climate series.

just two observed series, the CO2 and SOI. We use what is known as the Kalman smoother, applied to the VAR(2) model in state space form. To be precise, the information used to estimate the SST series and the error limits of this estimate, is just the pair of full CO2 and SOI series and the VAR(2) model (fitted to the four full series). The only information otherwise used from observations of the SST series is the linear trend by which it was corrected for VAR model estimation and which was used to restore that linear component after estimation. Neither was any information used from observations of the LST series. The upper plot in Figure 7 shows the observed SST series, its estimate and two standard error limits. Note that the estimates are not predictions from past values. Each value of the SST series is estimated from the whole sequence of the CO2 and SOI series. The lower plot in Figure 7 shows the difference between the observed and estimated series with the error limits.

We make the following remarks on these plots.

1. The estimated SST follows the pattern of the observed SST remarkably well, and even over periods where it is generally lower or higher, it follows the year to year movements well.

2. This similarity does not imply causality in either direction between the predicting series of CO2 and SOI and the predicted series of SST, because they may be related by mutual dependencies which we aim to model in later sections.

3. Though the observed SST varies by only a few tenths of a degree over the whole record, it is compiled from a large number of temperature measurements which are subject to much greater diurnal and annual variation. The precisely and objectively defined nature of the predicting series gives strong support to the claim that the well predicted observed temperature series is similarly well defined.

4. When the VAR model is fitted only to the observations of the four climate series before the year 2000, there is no visual difference in the plots of Figure 7 when they are constructed using this restricted model. In particular there is no visual difference in the SST estimated from the CO2 and SOI series over the years 2000 to 2014. This

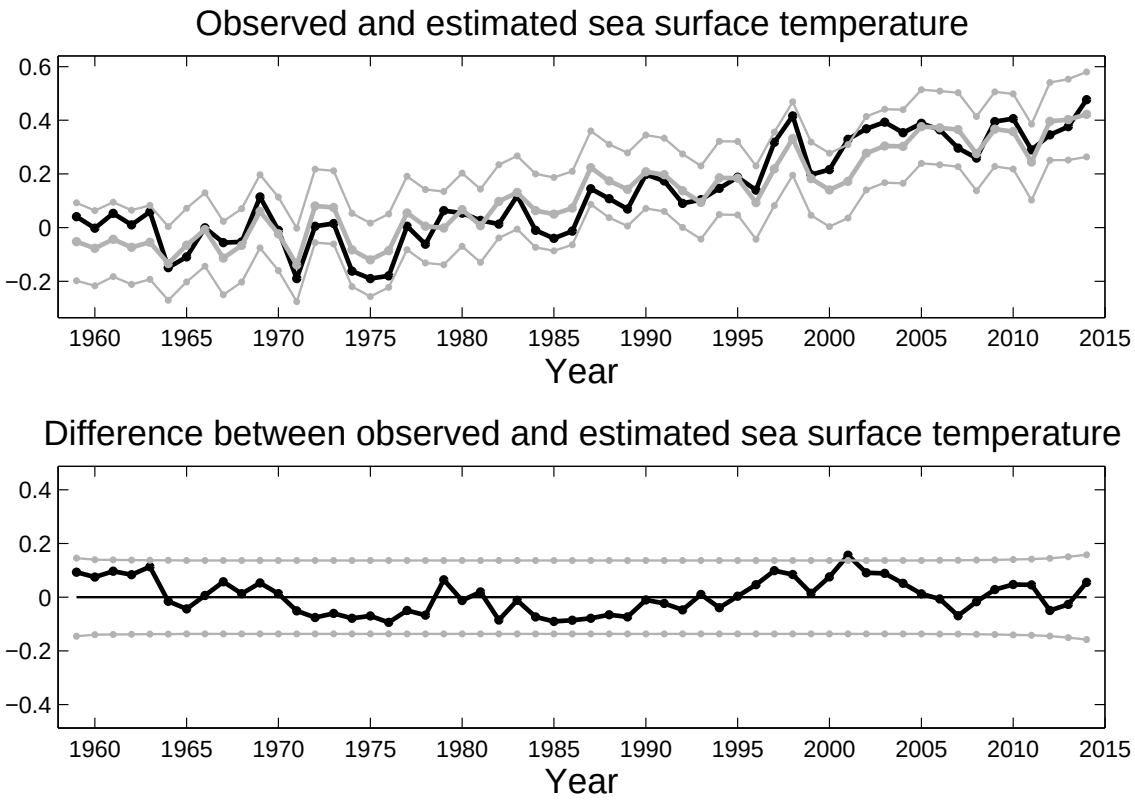

Figure 7: The upper plot shows the observed SST series (black line) and its estimate and two standard error limits (gray lines), derived from observations only of the CO2 and SOI series, using a VAR(2) model fitted to the four climate series. The lower plot shows the difference between the observed and estimated series with the error limits.

suggests that there is no essential change in the times series nature of the four climate series over this period.

5. We note from the plots in Figure 7 that the general level of the observed SST is above that of the estimated level over the period from 1995 to 2005. The apparent leveling off of the observed SST over the first decade of the new century is due to this surge, and its conclusion around the middle of that decade.

6. The difference series in the lower plot of Figure 7 has a typical red spectrum, rising at lower frequencies, and has no significant cyclical, or other statistical, features worthy of remark, though perhaps some climatic explanation may be found by inspection of its variation.

4 A structural VAR model for the climate series

By a structural VAR (SVAR) model we mean one in which each current value of the series may depend upon other current values besides a set of lagged values. This dependence

explains the correlation which would otherwise remain between the innovation series if dependence was allowed only upon past values, as in the standard VAR. The structural VAR model therefore results in structural innovations which are uncorrelated. The model equation is now

$$\Phi_0 x_t = \Phi_1 x_{t-1} + \Phi_2 x_{t-2} + \cdots + \Phi_p x_{t-p} + a_t, \tag{2}$$

where Φ_0 represents dependence between current values and the structural innovations a_t have diagonal variance matrix. Our aim is to identify and estimate this model with sparse forms of the coefficient matrices Φ_k which represent the structure of the dependence within and between the series.

We have previously propounded methods of identifying SVAR models, for example in Reale and Tunnicliffe Wilson (2001) and Tunnicliffe Wilson et al. (2015). This methodology is very much motivated by and based upon the graphical modeling procedures used to identify relationships between a general set of variables represented by directed acyclic graphs (DAGs). This is clearly set out in the books Whittaker (1990), Lauritzen (1996) and Edwards (2000). One of the main statistical tools used in this identification is the conditional independence graph (CIG) between the variables. A simple rule determines how the CIG for a set of variables may be derived from the DAG representing their dependence. A given CIG does not, however, necessarily determine uniquely the structure of this DAG. In many cases, however, a small number of possible DAG representations can be determined, one of which may be selected as the best, following their estimation by maximum likelihood.

The idea of acyclic dependence in a DAG, is that the variables may be ordered so that each is dependent only on a subset of those which are previous in the ordering. More traditionally, in econometrics, see Zellner and Theil (1962), a set of simultaneous equations, suggested by economic theory, may be used to represent the relationships between variables. Each equation may only involve a small subset of all the variables, upon which restrictions are imposed to ensure that the relationships may be uniquely identified and estimated. However, there is no requirement that the dependence be acyclic. These methods have been extended to represent simultaneous equation relationships between current values of a structural VAR model, see for example the much cited paper of Sims (1996).

In our earlier exposition, Reale and Tunnicliffe Wilson (2001), of graphical modeling as applied to identifying SVAR models, we restricted ourselves to acyclic relationships between current variables. However, in Tunnicliffe Wilson et al. (2015) we explored the possibility of using cyclic simultaneous equation representations, which we believe may be appropriate for the climate series which are the subject of this paper. Fortunately, the CIG between the variables, determined empirically by statistical analysis, may still be used in the identification of their dependence, whether or not this is acyclical.

Because our interest is restricted to linear relationships between variables described by their second order statistical moments, we can construct their CIG from their sample partial correlation graph (PCG). As described in Tunnicliffe Wilson et al. (2015), the PCG is formed from the data matrix $\mathbf{X}$ whose columns are the four mean and trend corrected climate series and their values to lag 2 - the order of their standard VAR representation. Their sample covariance matrix is then $\widehat{V} = \frac{1}{n}\mathbf{X}'\mathbf{X}$ and their sample inverse covariance matrix is computed

as $\widehat{W} = \widehat{V}^{-1}$. From its entries, $\widehat{W}_{i,j}$, the sample partial correlations can be calculated as

$$\widehat{\pi}(X_i, X_j) = \frac{-\widehat{W}_{i,j}}{\sqrt{\widehat{W}_{i,i}\widehat{W}_{j,j}}}. \tag{3}$$

A sample partial correlation between two variables is closely related to the t-value of the coefficient of one of them, in the regression of the other upon the whole set of variables. Under the hypothesis that $\pi(X_i, X_j) = 0$, the ratio

$$\frac{\widehat{\pi}(X_i, X_j)\sqrt{n-m+1}}{\sqrt{1-\widehat{\pi}(X_i, X_j)^2}} \tag{4}$$

is distributed as a t_{n-m+1} variable where $n-m+1$ is the number of degrees of freedom. We therefore reject the null hypothesis that $\pi(X_i, X_j) = 0$ at level α if

$$|\widehat{\pi}(X_i, X_j)| > \frac{t_{\alpha/2,n-m+1}}{\sqrt{t^2_{\alpha/2,n-m+1} + (n-m+1)}}. \tag{5}$$

where $t_{\alpha/2,n-m+1}$ is the corresponding critical value of the t_{n-m+1} distribution.

Our interest is purely in the relationships between current variables and between current and past variables, not between past variables. On applying this procedure to our climate variables we show in Figure 8 the CIG of their dependence. No link is shown between two variables for which the associated t-value defined in (4) is less than 1.645 in absolute value. The strength of a link is shown also by the style of line: either broken, thin solid or thick solid, if the t-value is greater than 1.645, 1.96 and 2.575, corresponding to test levels of 10%, 5% and 1%.

Our next step is to postulate an SVAR for which the corresponding CIG (or PCG) is given by that in Figure 8. We will also present this SVAR graphically using a similar diagram to that shown in the figure, but with directions attach to the links, indicated by arrow heads, showing which variables are used to explain the current variables at time t. All links from the past are naturally directed towards current variables. A link from one current variable to another specifies them as respectively explanatory and dependent variables. If the relationship so specified between the current variables is acyclic, then the SVAR may be estimated by OLS linear regression of each current variables on its specific set of explanatory variables. Otherwise, the Gaussian likelihood of the model is evaluated, by transforming it to the standard VAR form, and maximized numerically.

To illustrate the rule by which the CIG of an SVAR may be derived, we use that represented by the diagram in Figure 9.

The relationship between current variables in this figure is *not* acyclic. If it were we would call it a directed acyclic graph (DAG). As it is we note the cycle of links $\mathrm{CO2}_t \rightarrow \mathrm{SST}_t \rightarrow \mathrm{SOI}_t \rightarrow \mathrm{CO2}_t$, and we will call it a directed structural graph (DSG). However, for both types of graph the implied CIG (or PCG) between the variables can be derived using the moralization rules of Lauritzen and Spiegelhalter (1988). These are expressed in the language of graphical modeling in which each of the 12 variables shown in Figure 9 are referred to as nodes and for a given node its parent nodes are those from which it receives a directed link - i.e. its explanatory variables. The rule is then:

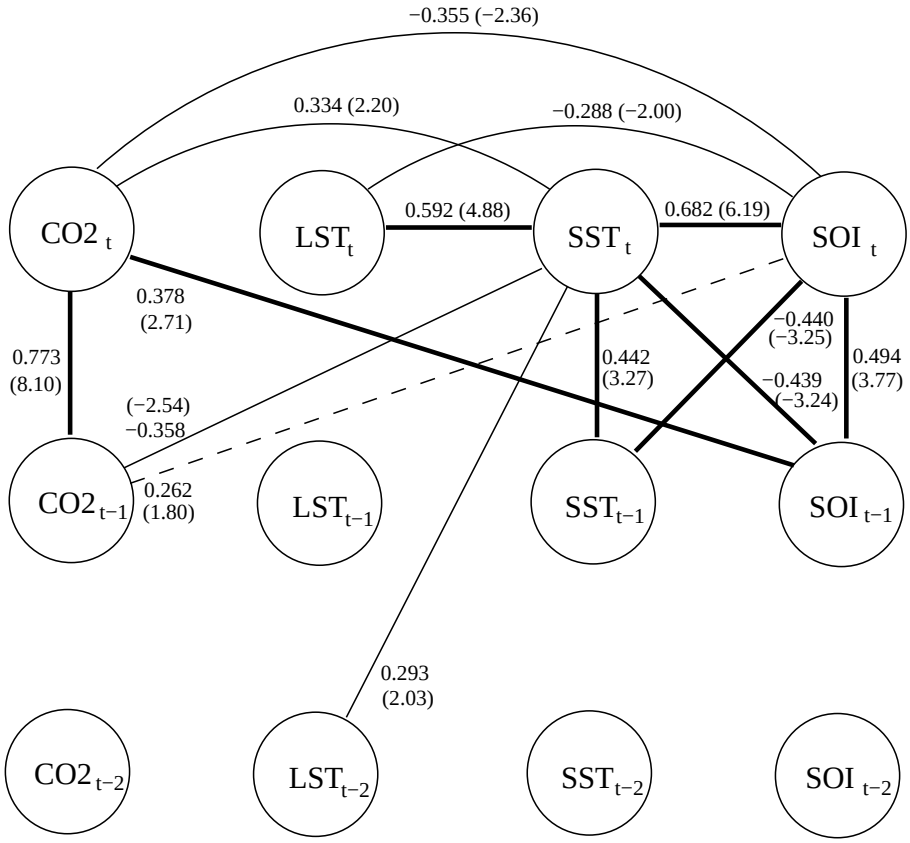

Figure 8: The sample CIG between the four climate series and their values to lag 2. Significant partial correlation between two variables is shown by a link, against which is shown the value of the partial correlation with its associated t-value in brackets.

1. For each node of the DAG or DSG insert an undirected edge between all pairs of its parent nodes, unless they are already linked by an edge. This is called marrying the parents (to make them moral).

2. Replace each directed edge in the DAG or DSG by an undirected edge.

Doing this for the graph in Figure 9 gives the graph in Figure 8, with a few extra links on which we comment shortly. Note that moralization in this case introduces the link $\mathrm{SOI}_{t-1} \to \mathrm{SST}_t$. Also the choice of direction for the link $\mathrm{SST}_t \to \mathrm{LST}_t$ avoids the introduction of a moralization link between SST_{t-1} and LST_t. Reversing the direction of the postulated link $\mathrm{SST}_t \to \mathrm{LST}_t$ would lead to the introduction of this moralization link, in conflict with the graph in Figure 8. Such considerations help to specify the postulated model. There are some other moralization links such as between SST_{t-1} and $\mathrm{CO2}_t$ that should be added. However, moralization links generally have a lower associated partial correlation and may not appear as significant in the PCG. One of the checks that we apply to a fitted SVAR is to compare all its implied partial correlations with the sample values used to form the PCG.

Of course, just because a given DAG or DSG (from here on we will just write DSG for

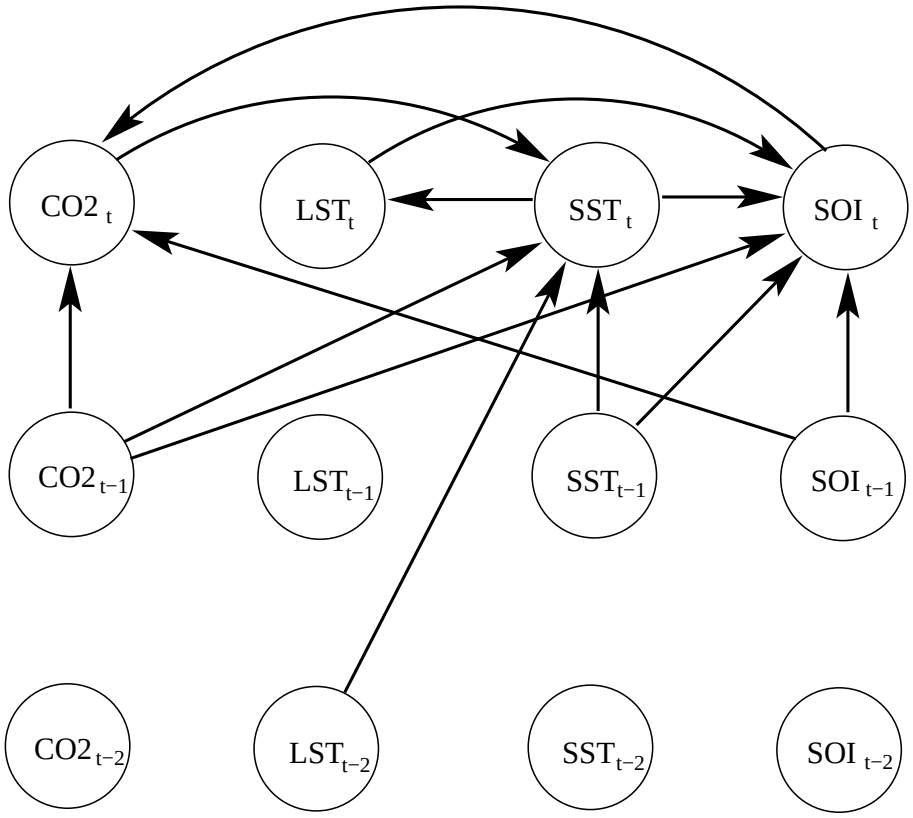

Figure 9: A diagram representing an SVAR which might explain the CIG shown in Figure 8.

this pair) is consistent with a given CIG does not mean that it is the correct model. For example a link in the CIG that might possibly be ascribed to moralization is not necessarily so explained. However, the true DSG will not contain links that are not present in the CIG, except by possible numerical coincidence, for example when a true link and a moralization link contribute canceling effects to give a zero partial correlation which would otherwise appear as a link in the CIG.

Identifying which of the links between current variables are not due to moralization, and the directions of the remaining links, is the key to identifying the DSG. All the links from past values that appear in the CIG can then be included in the DSG. Any due solely to moralization should be found, on fitting the model, to have coefficients that are not significant.

We have, in fact, for this example used a strategy which may be viewed as exploiting this last point. It is certainly not universally applicable, because it relies on the CIG having a high level of sparsity in its links. Even then it may not be successful because it relies on the fitting of what may be an over-specified model, for which there may be a range of likelihood equivalent parameter values, i.e. no unique set of estimates. We have, however, used it before with success in modeling term rates series in Tunnicliffe Wilson et al. (2015).

The strategy is to fit a DSG which includes every link in the CIG, with the directions from past to current naturally given by the arrow of time, but with every link between

current variables being made bidirectional. On fitting this model by Gaussian estimation the coefficients of low significance can be removed successively, with the level of significance being confirmed by differences in the Gaussian likelihood. This procedure was followed without difficulty, leading to the DSG identical to that shown in Figure 9 except for the bidirectional links $\mathrm{LST}_t \leftrightarrow \mathrm{SST}_t$. The link to the right, $\mathrm{LST}_t \rightarrow \mathrm{SST}_t$, has a low t-value of -1.43. However, this is based on a local quadratic approximation of the likelihood and on removing this term the deviance (minus twice the log likelihood) increases by 9.02. This suggests that both this term is significant and that the t-value approximation is poor in this case. We have however, on removing other terms in the model, found that their t-values were consistent with differences in deviance, and all the remaining terms are significant.

5 The final model, its interpretation and properties

The model so far identified is now subject to diagnostic checking of the residual series which in this case are the estimated structural model innovations. The largest sample cross-correlation between these at lag zero is 0.104, much less than the nominal two standard error limit of 0.267. However, there is a lagged cross-correlation with the value of -0.370 between the residuals of $\mathrm{CO2}_t$ and SST_{t-2}. We therefore introduced a corresponding further term into the model to explain this dependence and on estimation this term had a t value of -2.93. Although there remained several lagged cross-correlations on or just below their two standard error limits, this is to be expected among the 250 cross-correlations plotted up to lag 10 in Figure 11. The model with this extra term has a deviance in excess of that of the saturated standard VAR model by only 26.20, with 22 fewer coefficients, suggesting that we have not sacrificed any significant goodness of fit by using the sparsely parameterized structural model. The final model is displayed in Figure 10 with the coefficients displayed adjacent to the links and their t-values in brackets. The t-values are derived from local approximations of the deviance, except for the link $\mathrm{LST}_t \rightarrow \mathrm{SST}_t$ for which it is derived from the deviance difference.

Because of the cyclical nature of the relationship between current variables, we have to be careful how we interpret this diagram. It appears to present, for each current series value, its prediction given all the other current series values and the past values to lag 2. This is not true. If it were, this prediction would be the same as that of standard VAR model given the same series values. We will give further consideration of the properties of this model, but first present evidence that the model adequately represents the series.

The sample correlation matrix of the model residuals, or structural innovations, is:

$$\begin{pmatrix} 1.000 & & & \\ -0.089 & 1.000 & & \\ 0.044 & 0.009 & 1.000 & \\ -0.045 & -0.005 & 0.081 & 1.000 \end{pmatrix}. \tag{6}$$

All these entries are small and give no reason to doubt the assumption that the residuals are uncorrelated. The lagged cross-correlations of the residuals are shown in Figure 11.

Again, these generally lie within their bounds except that there appears to be some significant negative autocorrelation in SST at lag 4, which might be removed by a further

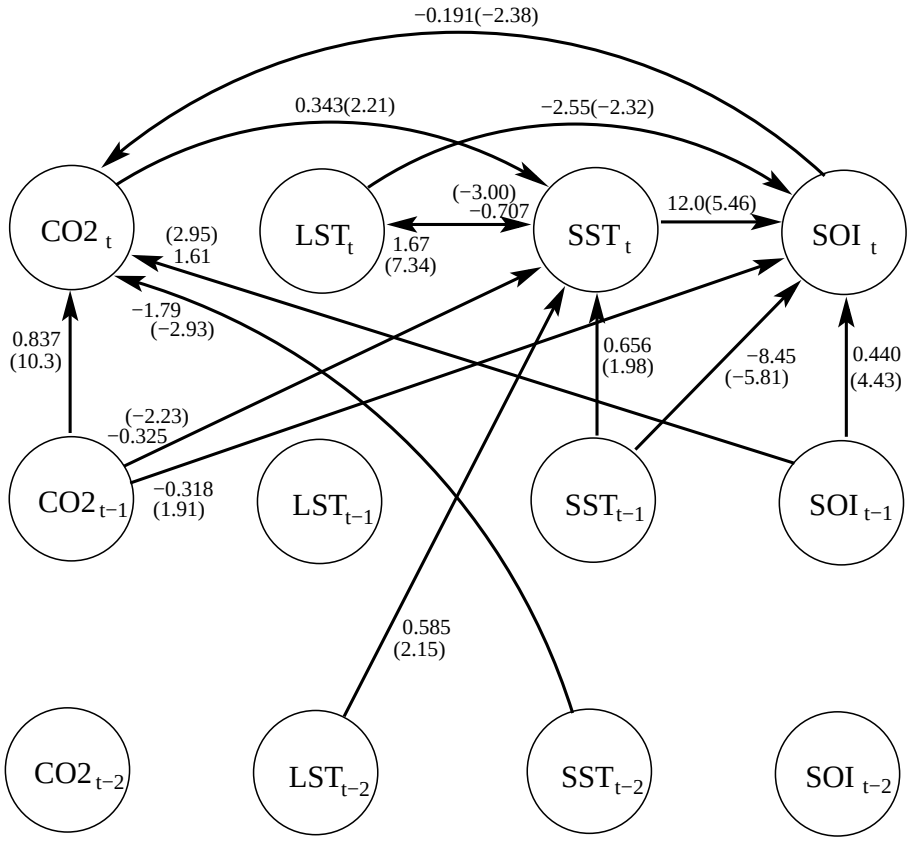

Figure 10: The final DSG fitted to the four climate series.

term. The overall measure of the magnitude of the lagged cross-correlations, the sum of squares of all 250 values scaled by the series length 56, is 130.52. This is a form of the multivariate portmanteau statistic, Hosking (1980), Li and McLeod (1981). It does not even exceed the expected value of 235 = 250-15 (the number of estimated coefficients) and again gives no evidence to doubt the model.

Finally, we compare the partial correlation properties of the fitted model with the sample values used to identify the model. These are compared in the plot of Figure 12.

There are 14 zero values implied by the model, as seen in the horizontal line of points in the center. The lowest of these has a corresponding sample value of -0.23 and the largest a sample value of 0.15, so lie within their two standard error limits of ±0.27. There are 23 non-zero model values arising from the 15 links in Figure 10 and their moralization links. There is a good correspondence between the model and sample values of these, subject to sampling fluctuation. Most of the sample partial correlations corresponding to moralization links are small and do not appear in Figure 8. Again, this plot generally supports the model.

It is appropriate here to refer back to the partial coherency graph of Figure 4. The model in Figure 10 implies that the partial coherency graph should be of the form shown in Figure 4, except for the addition of a link between LST and SOI. The main point is the separation between CO2 and LST which have no connecting links in Figure 10, or consequent moralization links. And there is no link between these series in Figure 4. The comparison

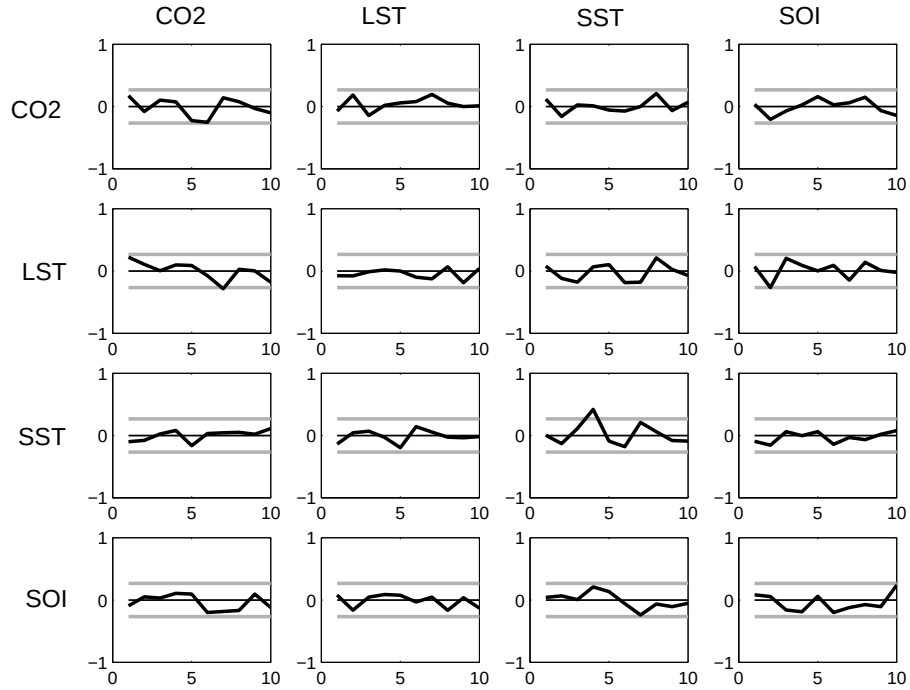

Figure 11: The cross-correlations up to lag 10 between the residual series of the DSG model. The gray bands show their approximate 2 standard error limits.

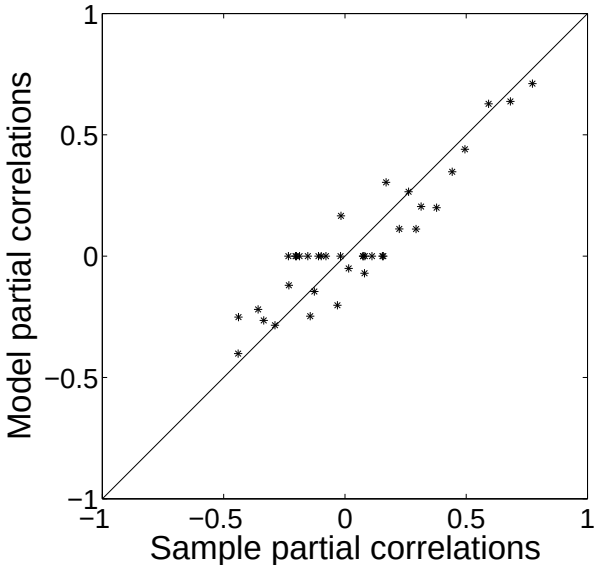

Figure 12: The model partial correlations up to lag 2 plotted against the corresponding sample partial correlations.

of these two graphs could highlight deficiencies in the SVAR, though not in this case. The partial coherency graph also requires little effort to compute and present.

We return to interpretation of the model. We have said that we cannot, for example,

take the predictor of SST_t in terms of all the other variables to be the linear combination of $\mathrm{CO2}_t$, LST_t, $\mathrm{CO2}_{t-1}$, SST_{t-1} and LST_{t-2} as indicated by its parents in Figure 10. However, we *could* do so if at any time the cyclical feedback links from SST_t to LST_t and SOI_t to $\mathrm{CO2}_t$ were broken. As it is we have to solve the simultaneous equations relating the contemporary variables to determine this predictor.

A widely used property of an econometric time series model is its impulse response function. This plots the future effect on the different series of a unit shock to the innovation in one of the series - in our case a structural innovation. Figure 13 shows the impulse response, and its cumulative value the step response, of the SST series, given a shock in CO2.

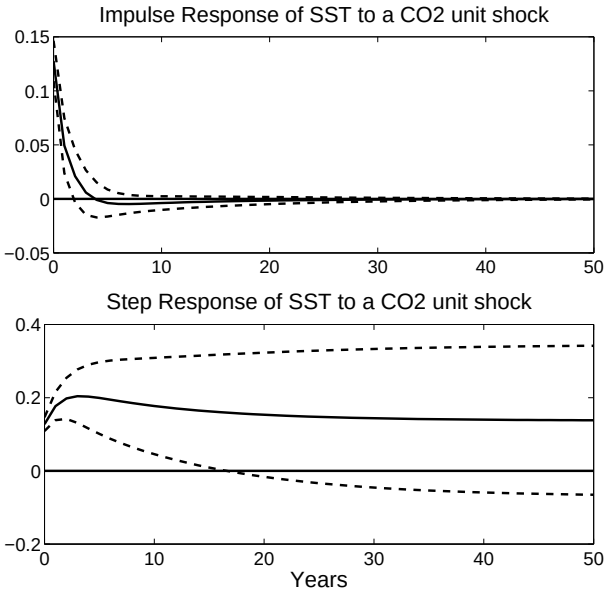

Figure 13: The impulse and step responses of SST to a unit shock in CO2, with one standard error limits.

One standard deviation error limits are shown, and even these are very wide for high lead times, in consequence of the fact that long-term information about the model has been removed by trend correction. Before commenting on these we show also, in Figure 14, the impulse and step responses of CO2 to a unit shock in CO2.

A unit innovation shock to CO2, even in the year zero at which it enters the system, leads to a less than unit increase in CO2 in that year. This is because of the negative feedback within that year from SOI to CO2, an effect which was previewed in the lowest plot of Figure 5. That plot shows the much larger positive response of CO2 occurring in the following year, which is supported in the final model of Figure 10 by the link $\mathrm{SOI}_{t-1} \rightarrow \mathrm{CO2}_t$. This positive feedback results in the step response to a series of unit shocks in CO2 leading to a continuing build up of CO2, and at the same time the greenhouse effect of the CO2 leads to a sea temperature increase. The model represents the effect that warm seas can absorb less CO2, leading to further net increase in the level of CO2.

Although the error limits on the responses are wide, the ratio of eventual level of the step

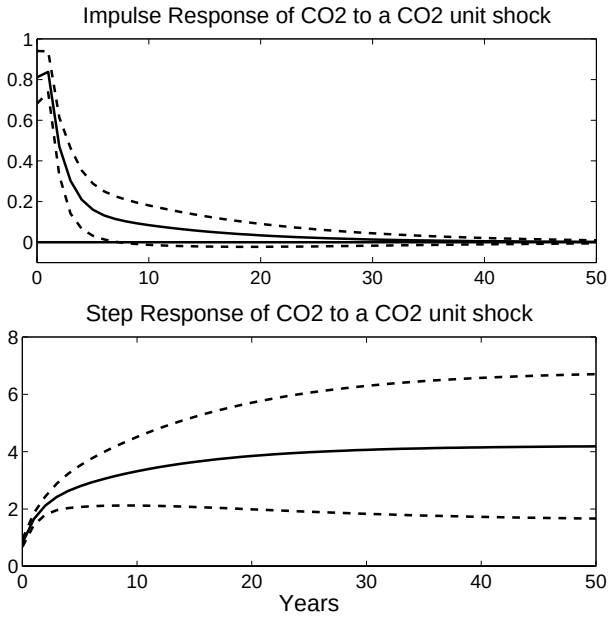

Figure 14: The impulse and step responses of CO2 to a unit shock in CO2, with one standard error limits.

response in SST to that of CO2 is approximately 0.033. The ratio in the trend slope of SST to that of CO2 seen over the period 1975 to 2005 where the slopes are steepest, in Figure 1, is somewhat less at 0.012. However, given the uncertainty in the model properties at higher lead times, these are not wildly inconsistent.

We also remark on the signs of some of the links between current variables. With our definition of the sign of SOI the current value is strongly positively dependent on the current value of SST, though the immediate past value of SST has a largely compensating negative effect - again as previewed in the upper plot of Figure 5. A substantial decrease in SST would have a negative effect on SOI and after a delay of one year this will led to an increase in CO2 associated with an El-Niño event.

Finally, the reciprocal roots of the SVAR model operator have maximum modulus 0.9133. This corresponds to a real root, with the next largest in absolute value being a pair of complex conjugate roots with modulus less than 0.5. This suggests that there may be a unit root process underlying the series, most likely deriving from the trend like behavior of CO2, with LST and SST being co-integrated with this.

6 Conclusion

Our analysis of these four climate series lead directly to a structural SVAR representation of the trend corrected series. The model is consistent with the trending appearance of the original CO2 and SST series and the terms in the model appear to have interpretation consistent with known climatological relationships.

References

H. Akaike. A new look at statistical model identification. *IEEE Transactions on Automatic Control*, AC–19(2):716–723, 1973.

R. Dahlhaus. Graphical interaction models for multivariate time series. *Metrika*, 51:157–172, 2000.

R. Dahlhaus and M. Eichler. Causality and graphical models in time series analysis. In P. J. Green, N. L. Hjort, and S. Richardson, editors, *Highly Structured Stochastic Systems*. Oxford: Oxford Universiity Press, 2003.

D. Edwards. *Introduction to graphical modelling*. New York: Springer-Verlag, 2000.

J. R. M. Hosking. The multivariate portmanteau statistic. *Journal of the American Statistical Association*, 75:602–608, 1980.

C. M. Hurvich and C.-L. Tsai. Regression and time series model selection in small samples. *Biometrika*, 76(2):292–307, 1989.

S. L. Lauritzen. *Graphical models*. Oxford: Oxford University Press, 1996.

S. L. Lauritzen and D. J. Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society: Series B*, 50:157–224, 1988.

W. K. Li and A. I. McLeod. Distribution of the residual autocorrelations in multivariate arms time series models. *Journal of the Royal Statistical Society: Series B*, 43:231–239, 1981.

H. Lütkepohl. *Introduction to multiple time series analysis*. New York: Springer-Verlag, 1993.

M. Reale and G. Tunnicliffe Wilson. Identification of vector AR models with recursive structural errors using conditional independence graphs. *Statistical methods and applications*, 10:49–65, 2001.

G. C. Reinsel. *Elements of multivariate time series analysis*. New York: Springer-Verlag, 1993.

C. A. Sims. Are forecasting models usable for policy analysis. *Federal reserve bank of Minneapolis quarterly review*, 10:2–16, 1996.

G. Tunnicliffe Wilson, M. Reale, and J. Haywood. *Models for dependent time series*. New York, CRC Press, 2015.

J. C. Whittaker. *Graphical models in applied multivariate statistics*. Chichester: Wiley, 1990.

P. C. Young. Hypothetico-inductive data-based mechanistic modelling, forecasting and control of global temperature. Technical report, Lancaster Environment Center, Lancaster University, UK, 2014. URL `http://captaintoolbox.co.uk/Captain_Toolbox.html/Captain_Toolbox.html`.

A. Zellner and H. Theil. Three-stage least squares: Simultaneous estimation of simultaneous equations. *Econometrica*, 30:54–78, 1962.