INTERNATIONAL JOURNAL OF CLIMATOLOGY
Int. J. Climatol. **20**: 1–18 (2000)

IS THE NORTH ATLANTIC OSCILLATION A RANDOM WALK?

DAVID B. STEPHENSON[a,*], VALENTINA PAVAN[b] and ROXANA BOJARIU[c]
[a] *Météo-France, 42 Avenue Gaspard Coriolis, 31057 Toulouse, France*
[b] *CINECA, via Magnanelli 6/3, 40033 Casalecchio di Reno, Italy*
[c] *NIMH, Sos. Bucuresti-Ploiesti no. 97, 71552 Bucharest, Romania*

Received 8 June 1998
Revised 22 March 1999
Accepted 9 May 1999

ABSTRACT

The North Atlantic Oscillation (NAO) is a major mode of large-scale climate variability which contains a broad spectrum of variations. There are substantial contributions from short-term 2–5 year variations, which have clearly marked teleconnections. Decadal trends are also apparent in the historical record of the NAO and may be due to either stochastic or deterministic processes. Evidence is presented that suggests the NAO exhibits 'long-range' dependence having winter values residually correlated over many years. Several simple stochastic models have been used to fit the NAO SLP (sea-level pressure) wintertime index over the period 1864–1998, and their performance at predicting the following year has been assessed. Long-range fractionally integrated noise provides a better fit than does either stationary red noise or a non-stationary random walk. Copyright © 2000 Royal Meteorological Society.

KEY WORDS: North Atlantic Oscillation; NAO; variability; stochastic processes; stationarity; trends; long-range dependence; forecasting; surface temperature

1. INTRODUCTION

Since at least the 18th century, it has been known in Denmark that when the winter there is severe, the winter in Greenland tends to be milder than normal, and vice versa (van Loon and Rogers, 1978). This remarkable opposition is part of a prominent planetary-scale mode of climate variability known as the 'North Atlantic Oscillation'. The North Atlantic Oscillation (NAO) is associated with a large-scale net displacement of air between the subtropical high near the Azores and the low pressure region near Iceland and the Arctic region (Teisserenc de Bort, 1883; Hildebrandsson, 1897; Walker, 1924). High index winters have stronger mean westerly flow over the North Atlantic and western Europe associated with a deeper than normal Icelandic low and a stronger than normal Azores high. Stronger westerly flow advects more warm maritime air over Europe and more cold Arctic air over Greenland and the northwest Atlantic, thereby giving rise to the anti-correlation between temperatures in Greenland and Denmark. The NAO has strong impacts on the climate and environment of both western Europe (Rogers, 1990; Hurrell, 1995) and eastern Europe (Bojariu and Cotariu, 1996). There is evidence that the NAO can influence wheat production in Northern Europe (Meinardus, 1900; Kettlewell *et al.*, 1999) and oceanic ecosystems in the North Atlantic ocean (Fromentin and Planque, 1996).

After more than a century of scientific investigation, the fundamental mechanisms determining the evolution of the NAO are still far from being elucidated. For example, it is not clear whether the NAO is simply due to the aggregation of stochastic weather events, or whether ocean dynamics in the North Atlantic ocean play an active role in controlling the evolution of the NAO. Improved understanding of causal mechanisms of the NAO may enable skilful climate forecasts to be made seasons to years in

* Correspondence to: University of Reading, Department of Meteorology, Earley Gate, PO Box 243, Reading, RG6 6BB, UK; e-mail: stephen@cict.fr

CCC 0899–8418/2000/010001–18$17.50

advance, with obvious socioeconomic benefits. It is also not clear whether possible anthropogenically-induced climate change has been responsible for the recent increasing trend in the NAO since the mid-1960s. More detailed understanding of previous NAO behaviour may ultimately allow us to conclusively discriminate between possible man-made changes and so-called natural variations in the NAO.

This study investigates various aspects of NAO variations present in the sea-level pressure and surface temperature observations briefly described in the Section 2. Section 3 discusses long-term variations and trends that have occurred in the NAO, and Section 4 explains how year-to-year differencing can be used to separate out such features. Section 5 then uses differences to extract short-term variations in the NAO. Section 6 proposes some simple stochastic models for describing the fluctuations in the NAO, and fits to these models are compared in the Section 7. Section 8 concludes the article with a summary and some speculations.

2. DATA USED IN THIS STUDY

The NAO has clear signatures in several different climatic variables such as temperature, pressure, rainfall, etc. This investigation will make use of wintertime means of observed sea-level pressure (SLP) and surface temperature. For both temperatures and pressures, the quoted year refers to the year of the spring months and not the year of the preceding December.

2.1. A SLP index of the NAO

The difference in SLP between the centres-of-action near to Iceland and the Azores has long been used to summarize the mean westerly flow over the North Atlantic region (Walker, 1924). Hurrell (1995) defined a winter NAO SLP index for 1964–1998 based on the difference between standardized December–March mean SLPs measured at Lisbon (Portugal) and Stykkisholmur/Reykjavik (Iceland).[1] It has been widely used in recent NAO studies such as Hurrell (1995, 1996), and Hurrell and van Loon (1997). Seasonal averaging is required in order to suppress the noisy intraseasonal fluctuations in monthly mean SLP. Figure 1(a) shows the evolution of this index from 1900 to 1994, with decadal trends depicted by 10-year running medians.

In addition to this station index, we also make use of globally gridded monthly mean SLPs for 1959–1994 obtained from the NCEP/NCAR 40-year reanalyses.[2]

2.2. Surface temperatures

Numerous studies have defined the NAO using surface temperatures rather than SLPs (e.g. Hann, 1890; Loewe, 1937; van Loon and Rogers, 1978). Temperatures have the advantage over pressures of having more persistence from 1 month to the next, and also have centres of action which are more geographically locked to fixed features such as continents. In this study, we have analysed the Jones *et al.* (1986) temperature data set which is freely available from http://www.cru.uea.ac.uk/. The data set was obtained by averaging observed air temperatures at 2 m above the ground and temperatures at the sea surface into 5° latitude–longitude boxes over the whole globe (Jones *et al.*, 1986).

Figure 1(b) shows wintertime mean temperatures derived from this data set by averaging over large regions covering northwest Europe (7.5°W–47.5°E, 52.5°N–72.5°N) and the northwest Atlantic (22.5°W–77.5°W, 52.5°N–72.5°N). There is a clear anti-correlation between the temperatures of the two regions as was realized by the Danes in the 18th century (van Loon and Rogers, 1978). Wintertime mean temperatures over NW Europe and NW Atlantic are (anti-) correlated with a product moment correlation of -0.50 over the period 1900–1994. NW Europe and NW Atlantic wintertime mean temperatures also have strong correlations of 0.60 and -0.59, respectively, with the December–March NAO SLP index over 1900–1994. When all the grid point values are used to estimate the area-weighted average over the northwest Atlantic region, the increasing number of colder observations that have become available close to Greenland gives the misleading impression that there has been a cooling trend of 5°C since 1900. To obtain reliable estimates of area-averages, we have only considered data after 1900, and have not included

ANY values from grid points that have less than 80 winter means defined out of the 95 winters from 1900 to 1994. This simple quality control approach is an effective way of dealing with the difficulties caused by taking area averages when there are a varying number of missing values. In this study, wintertime mean temperatures are obtained by taking the average of the monthly means at each grid point for December, January, and February (DJF).

3. LONG-TERM NAO VARIATIONS

3.1. Decadal trends

Decadal variations are apparent in the 10-year running medians of the NAO SLP index shown in Figure 1. For example, the increasing trend that started in the mid-1960s, which has provoked speculation about possible anthropogenic-induced climate change (Trenberth, 1990; Hurrell, 1996; Wallace *et al.*, 1996). Prior to this recent bullish episode, there was an equally strong decreasing trend from the late-1940s to the mid-1960s (Rogers, 1984), which is not easily explainable as a consequence of global warming. The decadal trends, however, may be caused by other deterministic influences such as ocean dynamics (Deser and Blackmon, 1993; Kushnir, 1994), or may simply be the result of aggregating values in noisy time series (Wunsch, 1999).

The presence of decadal trends can be assessed using the Kruskal–Wallis rank statistic to test the null hypothesis that all the decadal means come from the same underlying population against the alternative hypothesis that the decadal means come from different populations (Hollander and Wolfe, 1973). Applying this test to the 12 decadal means of the NAO SLP index calculated from 1870 to 1989, gives a Kruskal–Wallis H-statistic of 18.422. Under the two-sided null hypothesis, this statistic is expected to be χ^2 distributed with 11 degrees of freedom and so has a probability of 0.07 of occurring. Therefore, at 90%

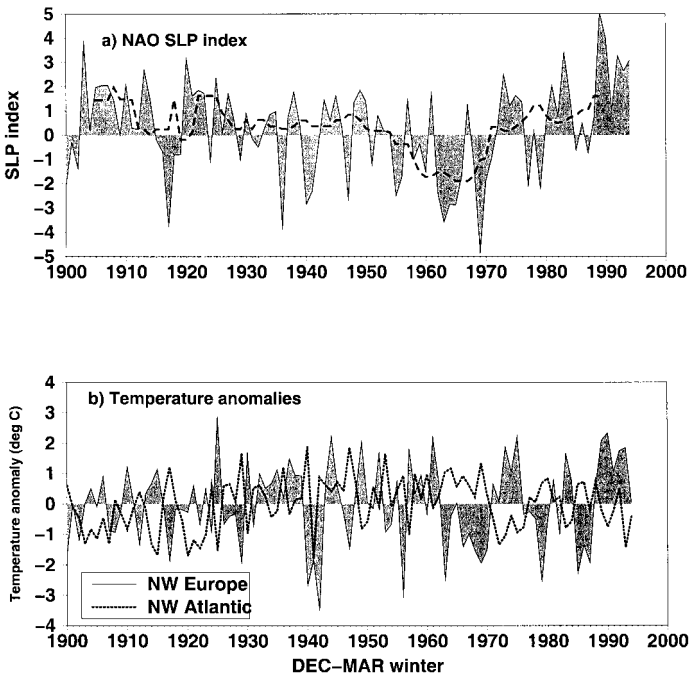

Figure 1. Evolution from 1900 to 1994 of: (a) the NAO SLP index—difference in standardized December–March mean SLPs between Lisbon (Portugal) and Stykkisholmur (Iceland). Decadal trends are indicated by the 10-year running median (dashed line); (b) December–February mean surface temperatures averaged over large regions covering NW Europe (7.5°W–47.5°E, 52.5°N–72.5°N) and the NW Atlantic (22.5°W–77.5°W, 52.5°N–72.5°N)

confidence, the null hypothesis that the decadal means come from the same population can be rejected. In other words, the NAO has stronger decadal variations than are likely to be obtained from sampling a white noise process. The main contributions are from the high index decades of 1900–1929, 1980–1989, and the low index decades of 1870–1989, and 1950–1969.

3.2. Long-range dependence

Many naturally occurring processes have values which remain residually correlated with one another even after many years. Such 'long-range dependence' has been noted in many diverse environmental quantities such as the Nile river minima (Hurst, 1951), global mean temperatures (Bloomfield, 1992), daily Italian rainfall amounts (Montanari *et al.*, 1996), and wind power variations from 1961 to 1978 over Ireland (Haslett and Raftery, 1989). It is worth noting that all these quantities are influenced to varying extents by the NAO, especially wind power over Ireland. 'Long-range dependence' does not signify the existence of a 'memory' having a particular time-scale, but implies instead that the process forgets slowly its past behaviour—autocorrelations decay hyperbolically in time rather than exponentially. Long-range dependence is reviewed in Beran (1994).

In Figure 2(a), it can be seen that despite being small, autocorrelations of the NAO SLP index retain small magnitudes even up to large lags. This behaviour differs from the fast exponential decay expected for short-range processes such as red noise. Evidence for long-range dependence is also provided by the increasing power at low frequencies in the power spectrum of the NAO SLP index (Figure 2b). For short-range processes, the spectral power density asymptotes to a constant value at low frequency as can be seen in the spectrum for the red noise fit to the NAO SLP index shown in Figure 2(b). Such flattening out does not occur for the spectrum of the NAO SLP index shown in Figure 2(b), which instead appears to have a singularity (pole) at zero frequency. Wunsch (1999) has estimated that the NAO SLP index has a spectral power density that is closely approximated by $P(\omega) \approx 0.66\omega^{-0.22}$, where ω is the angular frequency. The singular increase of power as $\omega \to 0$ is a characteristic feature of long-range dependence (Beran, 1994).

3.3. Aggregated variance

The variance of sample means of n independent values satisfies the well-known expression $\mathrm{var}(\bar{z}_n) = \sigma^2 n^{-1}$. This expression is also expected to hold for series generated by short-range processes, with values that become independent of one another after long enough periods. Figure 2(c) shows the 'aggregated variance' $\mathrm{var}(\bar{z}_n)$ estimated for consecutive sample means of n winters of the standardized NAO SLP time series from 1864 to 1998. The variances of the means all exceed the n^{-1} value expected for short-range time series that become uncorrelated after long enough times. For example, the variance of decadal means is equal to 0.187, which is almost twice the value of 0.1 expected from the n^{-1} law.

The linearity in the log–log plot in Figure 2(d) suggests the presence of a power law scaling regime having $\mathrm{var}(\bar{z}_n) = \sigma^2 n^{-\alpha}$. From a weighted least squares linear fit to the points in Figure 2(d), an estimate of $\alpha = -0.79$ is obtained for the NAO SLP index scaling exponent. This estimate is almost two standard errors greater than the value of $\alpha = -1$ expected for a short-range white noise process of the same length. The standard error in α is estimated to be 0.11 from aggregated variance fits to 1000 Monte-Carlo realizations of 135-year-long series of white noise.

4. DETRENDING BY USING DIFFERENCES

As was noted by Helland-Hansen and Nansen (1920), North Atlantic variations often exhibit substantial short-term year-to-year fluctuations superimposed on longer term trends. The NAO is no exception and contains spectral power over a broad band of different frequencies, with no significant preference for any particular frequency band (Wunsch, 1999). Decadal trends, which typically explain less than 10% of the total variance, can nevertheless confound analyses of short-term variability and predictability. Because of

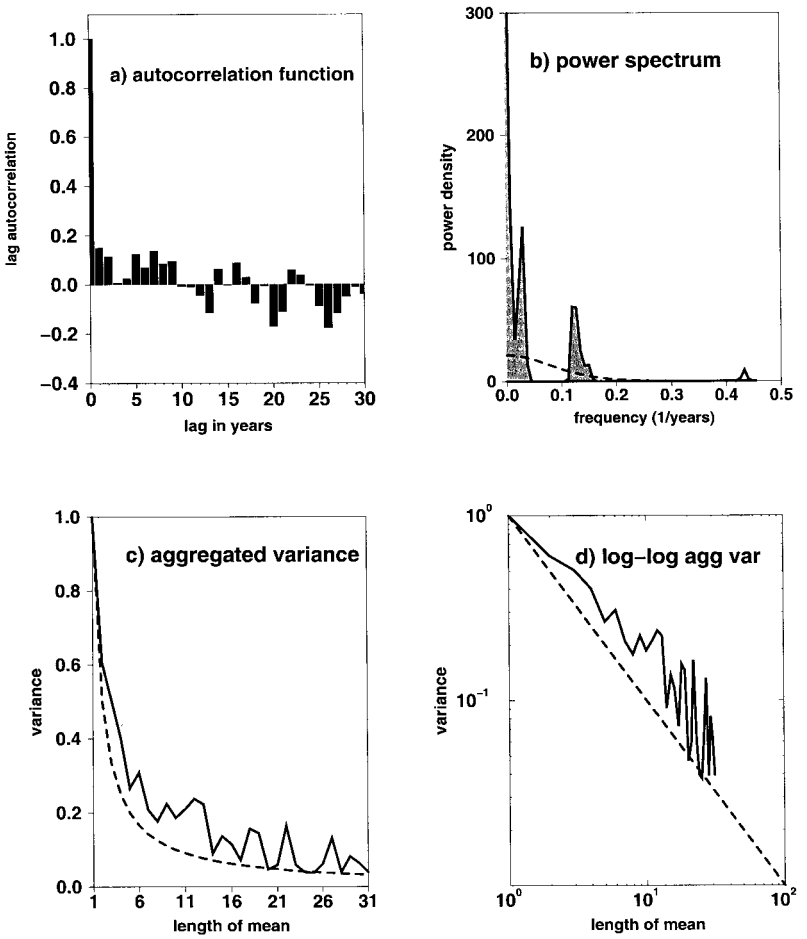

Figure 2. Evidence for long-range dependence in the NAO SLP index: (a) autocorrelation function showing slow decay of autocorrelations with lag; (b) spectral power density showing increasing power at low frequencies compared to the power density of a short-range AR(1) fit (dashed); (c) variance of means of increasing length compared to $1/n$ behaviour expected for independent values (dashed); (d) log–log plot of aggregated variance showing linear slope shallower than the slope of -1 expected for independent values (dashed)

increased noisiness due to atmospheric instabilities, great care needs to be exercised when analysing and forecasting variations in mid-latitudes.

4.1. Year-to-year variations

It took the genius of Bjerknes (1962) to realize that long-term decadal variations in the North Atlantic ocean may have quite different natures and causes to shorter term variations. To isolate short-term variations, Bjerknes and his students examined many maps of year-to-year differences in North Atlantic SSTs and SLPs for the period 1880–1915 (Bjerknes, 1964). Taking the difference between the value in one year and the previous year, $\Delta z_y = z_y - z_{y-1}$, is a standard and widely-used method for detrending time series (Box and Jenkins, 1976). It has the advantage over removing a linear fit in that it can also remove piecewise linear trends by converting them into constant terms. It also has the advantage over methods such as Fourier filtering or removing linear fits in that it is local in time. This is a desirable property since trends often behave differently at different times. For example, the NAO SLP index had a decreasing decadal trend from 1940 to 1960, but then an increasing one from 1960 to 1990. Any time series can be decomposed into the sum $z_y = \bar{z}_y + z'_y$, of the 2-year moving average $\bar{z}_y = (z_y + z_{y-1})/2$ and the high-pass

residual $z'_y = (\Delta z_y)/2$. The high-pass filter $\Delta/2$ attenuates the amplitude of low-frequency signals by a factor of $\sin(\omega\tau/2)$ where ω is the angular frequency and τ is the time between samples (e.g. 1 year). Therefore, biennial signals with periods of 2 years suffer no attenuation whereas signals with periods of 4 years have their amplitudes attenuated by a factor of $\sqrt{2}$. Decadal and lower frequency variability are attenuated to less than 30% of their original amplitudes. This justifies Bjerknes (1964) method of using year-to-year differences for extracting the '2–5 year short-term trends' (his words). Year-to-year differences in budget quantities such as the total heat content of the upper ocean can also be physically interpreted as being the total amount of heat received throughout the whole of the preceding year.

4.2. Measuring roughness

The ratio of the variances, $\mathrm{var}(\Delta z)/\mathrm{var}(z)$, can be used to measure the 'roughness' of a time series and is known as either the von Neumann ratio (von Neumann, 1941, 1942).[3] The ratio is identical to $2(1-\phi_1)$ where ϕ_1 is the moment estimate of the lag-1 autocorrelation.

Figure 3 shows the geographical distribution of 50 times the square root of the von Neumann ratio, in other words, the percentage of the total standard deviation (S.D.) explained by short-term variations. Short-term variations contribute to more than 50% of the total root mean square amplitude over all the Northern Hemisphere. The ratio is larger over continental areas than over oceans due predominantly to the reduced thermal inertia over the drier land surface. A notable exception is the local maximum over the eastern equatorial Pacific caused by short-term El Niño variations. There is a marked minimum south of Greenland related to deep vertical mixing, and also intriguing local minima in the Sea of Okhotsk, and in the Pacific ocean northeast of Japan.

5. SHORT-TERM VARIATIONS

This section will focus in more detail on the short-term variations obtained by taking year-to-year differences of the data. Analysis of short-term variations has the advantage that the effective sample size is larger than for data contaminated with decadal trends. There is less chance of obtaining spurious correlations and so more robust estimates can be made of spatial patterns of behaviour. Since differencing attenuates the amplitude of dominant ENSO variations by about half, and also removes decadal trends, it can provide a very useful technique for investigating short-term climate variability.

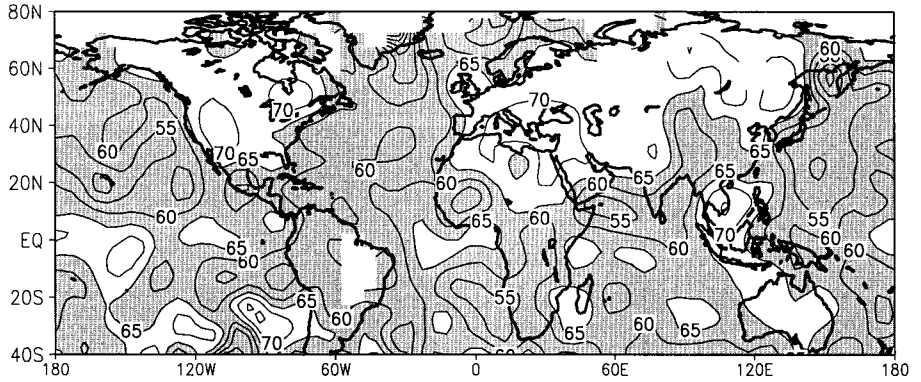

Figure 3. Relative contribution of short-term variations to total root mean square variation of December–March mean surface temperatures (in %). Uncorrelated white noise variations would give a ratio 71%. Shaded regions with less than 65% have less short-term contributions and more power at lower frequencies

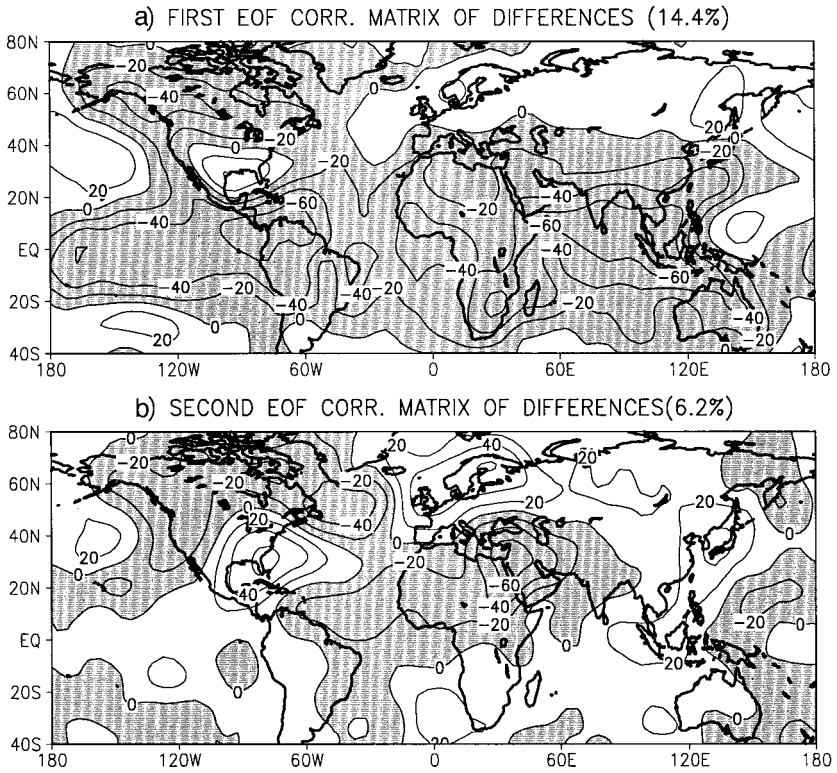

Figure 4. Leading eigenvectors of the correlation matrix of year-to-year differences in gridded winter (DJF) surface temperatures from 40°S–80°N: (a) first leading eigenvector showing strong correlations with the tropics; (b) second leading eigenvector showing a quadrupole pattern over the North Atlantic and surrounding land regions, containing the well-known NAO dipole between Greenland and NW Europe. Eigenvectors are normalized so as to represent correlations in %.

5.1. Spatial correlations: teleconnections

To summarize the correlations between the short-term variations in climate of different regions, we have performed a principal component analysis of the differenced surface temperature wintertime means from 1900 to 1994 over the longitudes 40°S–70°N. Since land temperature anomalies are typically five to ten times larger than anomalies over sea, the correlation matrix must be used instead of the covariance matrix in order to extract the main factors which contribute to correlations between short-term variations in temperatures in different regions. The resulting leading eigenvectors are shown in Figure 4 and explain 14.4% and 6.2%, respectively, of the total correlation.[4] Tests reveal that both these modes are robust features that are not unduly sensitive to either the choice of the spatial domain or the time period. The eigenvectors have been normalized so that they can be interpreted directly as maps of correlations that would have been obtained by correlating the data with the respective principal components.

The leading eigenvector in Figure 4(a) has strong zonally coherent tropical correlations particularly prevalent in the Indian ocean sector, with also some indication of a Pacific North America wave pattern. The leading principal component is anti-correlated with the backward difference of the June–September all-India rainfall monsoon index ($r = -0.47$), which strongly suggests that this mode is of tropical origin and is related to the tropical biennial oscillation (Meehl, 1997).

The second leading principal component is strongly correlated with the differenced NAO SLP index ($r = 0.66$) and the associated eigenvector in Figure 4(b) shows the characteristic NAO temperature dipole expected between NW Europe and the NW Atlantic. Interestingly, there are also strong correlations with temperatures over the southeastern USA and the Middle East which were also briefly discussed by Walker (1924). The NAO temperature signature is a quadrupole pattern consisting of the well-known northern

dipole together with an opposite phase dipole in the subtropics. Over the North Atlantic ocean, the two leading modes of short-term temperature variability both resemble the short-term pattern discovered by Bjerknes (1962, 1964). However, the leading eigenvectors are distinguished from one another on a global scale by the presence and absence of correlations with the tropics. The global nature of the NAO is confirmed by its emergence as the second leading eigenvector of the correlation matrix of near-global short-term temperature variations.

5.2. Covariance of SLP and temperature

To briefly summarize the correlations between short-term variations in surface temperature and tropospheric flow, we have also performed a canonical correlation analysis of the year-to-year differences in winter temperatures (DJF) and year-to-year differences in NCEP/NCAR reanalysis wintertime (DJF) SLPs over the period 1959–1994. The differencing helps to eliminate decadal trends that can sometimes give rise to spurious correlations. In this study, the temperature and SLP data are prefiltered by first projecting onto the first five leading principal components. The first two leading eigenvectors of the differenced wintertime (DJF) SLP covariance are shown in Figure 5(a) and (b) and explain 32% and 12% of the total variance. The first eigenvector projects strongly on the Iceland and Azores centres-of-action of the NAO, yet also has strong correlations extending zonally out of the Atlantic sector that encircle the Arctic region (Thompson and Wallace, 1998). The second eigenvector represents more a transpolar wavetrain pattern connecting the Pacific and Atlantic sectors. Figure 5(c) and (d) show the SLP and temperature patterns that give maximum canonical correlation ($r = 0.94$). The SLP pattern has a strong projection on the North Atlantic Iceland–Azores dipole, yet does not resemble the leading SLP eigenvector in the Pacific sector. The pattern is a strong mixture of the two leading eigenvectors and does not resemble the Arctic oscillation discussed in Thompson and Wallace (1998). The surface temperature pattern in Figure 5(d) resembles that of the second leading temperature eigenvector previously shown in Figure 4(b). It should be noted that canonical correlation analysis can not be used to infer causality. However, more in-depth studies suggest that the surface temperature pattern is primarily a response to atmospheric flow dynamics and not vice versa (Thompson and Wallace, 1998).

5.3. 2–3 year fluctuations

A substantial contribution to short-term variability comes from variations having periods between 2 and 3 years, as can be seen in the increased power spectral density in Figure 2(b). Such variations were particularly prevalent in the 1950s, and have been a robust feature of North Atlantic variability since at least 1701 (Deser and Blackmon, 1993; Cook *et al.*, 1997; Hurrell and van Loon, 1997; Tourre *et al.*, 1999). Quasi-biennial signals phase-locked with the annual cycle have been found in mid-latitude surface temperatures (Clayton, 1885; Landsberg *et al.*, 1963; Lamb, 1972), and are present in mid-latitude tropospheric winds and SLP variations (Angell and Korshover, 1963; Angell *et al.*, 1969; Ebdon, 1975; Trenberth, 1980; Trenberth and Shin, 1984). The cause of these short-term variations is still unknown but they could be due to either annual modulation of coupled atmosphere–surface feedbacks, or could be a subharmonic of the annual cycle generated by mid-latitude non-linearities. There has been speculation that the quasi-biennial variations may offer some hope of predictability in mid-latitudes (Brier, 1968; Madden, 1976). These tropospheric fluctuations do not appear to be strongly associated with the stratospheric quasi-biennial oscillation (Trenberth and Shin, 1984).

6. STOCHASTIC MODELS OF THE NAO

This section will consider some simple stochastic models that may be suitable for capturing some of the salient features of the NAO. In contrast to complex physically-based general circulation models, such statistical models are easy and quick to use and allow one to make many different forecasting tests. Furthermore, such models can provide invaluable and indispensable 'benchmark' control forecasts

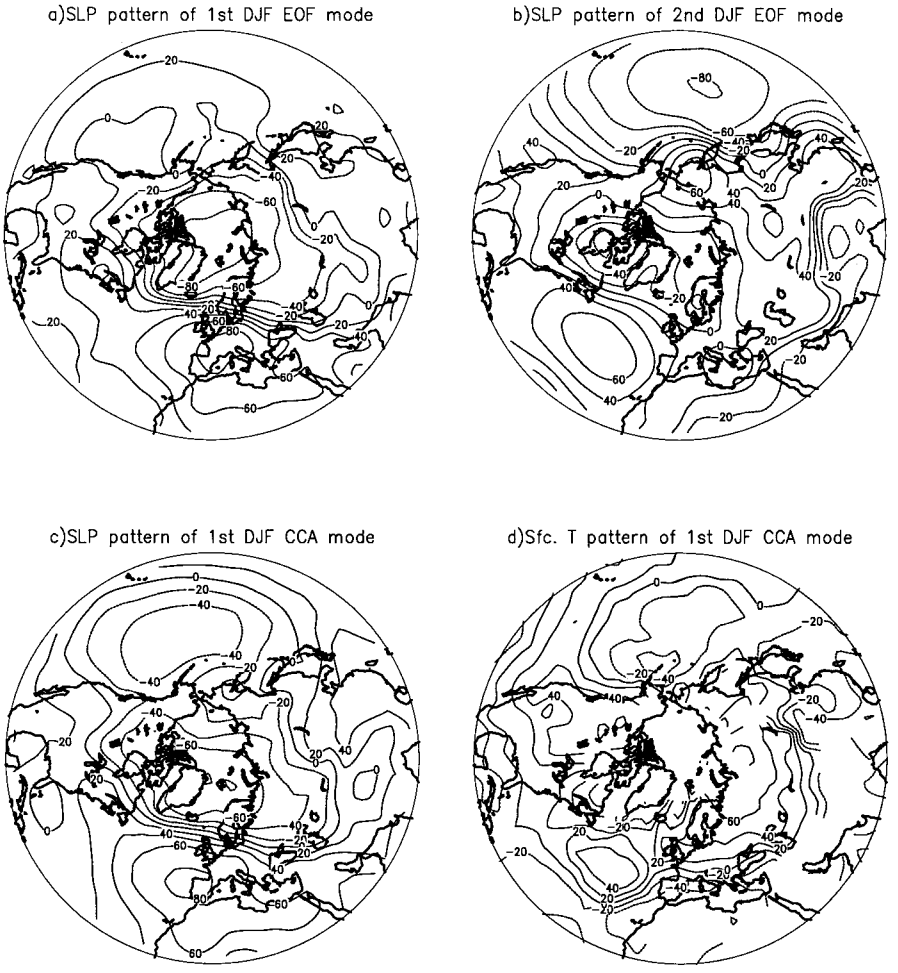

Figure 5. Short-term variations in tropospheric dynamics: (a) first eigenvector of covariance matrix of year-to-year differences in SLP 1959–1994 (32% of total variance) showing annular pattern with strong projection on the North Atlantic sector; (b) second eigenvector of covariance matrix of year-to-year differences in SLP 1959–1994 (12% of total variance) showing a transpolar wavetrain pattern; (c) pattern of differenced SLPs that has maximum canonical correlation with differenced surface temperatures; (d) pattern of differenced surface temperatures that has maximum canonical correlation with differenced SLPs. Patterns are normalized so as to represent correlations in %

necessary for assessing the possible forecast skill of more complex models. This section will describe four different time series models that have been used to fit the NAO SLP index. The estimated model parameters for the NAO SLP index from 1864 to 1998 are given in Table I. The quality of the fits will be examined in detail in Section 7.

Table I. Autoregressive time series models investigated in this study

Model	p	d	$\hat{\phi}_1$	$\sigma(\hat{\phi}_1)$
AR(1)	1	0	0.15	0.09
AR(10)	10	0	0.13	0.09
FAR(1)	1	0.13	0.005	–
RW	1	1	−0.48	0.08

All models are of the form $\Phi_p(B)(1-B)^d z_y = a_y$ where B is the backward shift operator $Bz_y = z_{y-1}$ and a_y is white noise (cf. Appendix A). Model parameters d and ϕ_k are estimated using maximum likelihood fits to the standardized NAO SLP index over the 135 winters from 1864 to 1998.

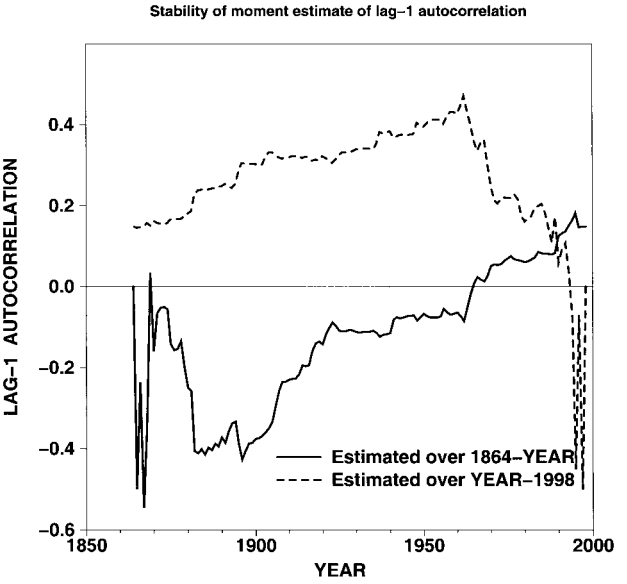

Figure 6. Estimates of the lag-1 autocorrelation coefficient for the NAO SLP index made over the period 1864-YEAR (solid line), and YEAR-1998 (dashed line). Note the absence of convergence in both cases even after 135 years

6.1. A stationary 'red noise' model: AR(1)

The 'red noise' paradigm of climate variability considers extra-tropical surface temperature anomalies to be the low-pass filtered response to stochastic heat fluxes associated with passing weather events (Davis, 1976; Hasselmann, 1976; Frankignoul and Hasselmann, 1977). The thermal capacity of the sea or land surface is responsible for generating serial correlations in surface temperature when forced with uncorrelated heat fluxes. This paradigm has been useful for understanding model generated sea surface temperatures (Griffies and Bryan, 1997; Saravanan and McWilliams, 1997). Incidentally, the word 'noise' has an oceanographic origin since it is derived from the Latin word 'nauseam', which describes the possible sea-sickness that can occur when sailing over rough (noisy) seas. The continuous time model proposed by Hasselmann (1976) is not directly applicable to discrete time series such as the wintertime NAO SLP index. However, a discrete analogue having the same low frequency limit is provided by the first order AR(1) autoregressive model:

$$z_y = \phi_1 z_{y-1} + a_y, \tag{1}$$

where z_y is the anomaly from the time mean for year y, and a_y is a Gaussian distributed innovation (shock). The maximum likelihood estimate of the lag-1 autocorrelation ϕ_1 for the NAO SLP index from 1864 to 1998 is quite small (0.15) yet is almost two standard errors larger than zero (Table I). Furthermore, due to the presence of substantial sampling fluctuations, the lag-1 estimate does not appear to have converged even with 135 years of data (Figure 6). The AR(1) model is one of the simplest time series models that can be used to model and forecast stationary time series.

6.2. A higher order autoregressive model: AR(10)

It can also be useful to consider higher order autoregressive processes such as:

$$z_y = \sum_{k=1}^{p} \phi_k z_{y-k} + a_y. \tag{2}$$

These models are special cases of the general class of Auto-Regressive Integrated Moving Average (ARIMA) models described in more detail in Appendix A. ARIMA(p, 0, 0) models have more parameters

$\{\phi_1, \phi_2, \ldots, \phi_p\}$, and can therefore sometimes provide better fits to a time series than do simple AR(1) models. For example, models of order $p \geq 2$ can capture oscillatory behaviour, whereas AR(1) processes always damp out exponentially. One disadvantage of higher order models is that extra uncertainty is introduced by the need to estimate more parameters. It is therefore advisable when making forecasts, to use models that do not have more parameters than necessary. This principle of parsimony is explained in more detail in Box and Jenkins (1976). In this study, we have considered the AR(10) model which uses the preceding ten winters to forecast the following winter.

6.3. A 'long-range' fractional differenced model: FAR(1)

Evidence was presented in Section 3 which suggests that the NAO exhibits long-range dependence. Stationary 'long-range' processes can be conveniently modelled using ARIMA models having fractional powers d of the difference operator (Hosking, 1981). To test this approach, we have examined the behaviour of the fractional differenced first order autoregressive model FAR(1) = ARIMA(1, d, 0) defined as:

$$\Delta^d z_y = \phi_1 \Delta^d z_{y-1} + a_y. \tag{3}$$

The fractional power of the difference operator is obtained using the binomial expansion:

$$(1-B)^d = 1 - dB + \frac{d(d-1)}{2!} B^2 - \frac{d(d-1)(d-2)}{3!} B^3 + \cdots, \tag{4}$$

which converges for values of d between -0.5 and 0.5. Non-stationary processes are obtained when $|d| > 0.5$ (Hosking, 1981). Such models have been successfully used to model and forecast economic time series (Geweke and Porter-Hudak, 1983), wind power in Ireland (Haslett and Raftery, 1989), and trends in global warming (Bloomfield, 1992). Maximum likelihood estimates of d and ϕ_1 given in Table I, are small and suggest once again that the NAO SLP index is close to being a stationary white noise process. Although ϕ_1 is estimated to be very close to zero, the estimate of the difference power d is positive and non-zero (0.15) suggesting a stationary long-range process. This estimate of d predicts values for the scaling exponent $\alpha = 2d - 1 = -0.7$ and power spectrum scaling $P(\omega) \approx \omega^{-2d} = \omega^{-0.3}$ in agreement with those obtained in the Section 3. The small estimate for d implies that the NAO SLP index is a stationary process having only a small amount of long-range dependence. Such a model is expected to explain only a small fraction of the total variance of the noisy time series (Hosking, 1981).

6.4. A non-stationary 'random walk' model: RW

Year-to-year differences can also be modelled using an AR(1) model:

$$\Delta z_y = \phi_1 \Delta z_{y-1} + a_y. \tag{5}$$

The values $\{z_y\}$ generated by this ARIMA(1, 1, 0) model (cf. Appendix A) are no longer tied to zero as they were in the AR(1) and AR(10) models, and can drift to arbitrarily large or small values. The variance of this non-stationary process increases indefinitely with time. When $\phi_1 = 0$, the model describes a pure random walk process in which z_y is incremented each year by independent Gaussian steps. The maximum likelihood estimate of ϕ_1 is given in Table I, and is close to the theoretical value of -0.5 expected for a differenced white noise process. Non-stationary models such as ARIMA(1, 1, 0) are widely used in economics and financial applications for describing and forecasting series having marked trends.

6.5. A comparison of some 500-year simulations

Figure 7 shows some 500-year-long realizations that have been simulated using the four different stochastic models. Decadal trends and periodicities are apparent in the running medians of all the simulated series, as were previously noted in the NAO SLP series. The decadal variations in the simulated series are the result of filtering random processes and have no underlying dynamical explanation. The

presence of apparent trends and periodicities in filtered random processes is known as the 'Slutsky–Yule effect', and can easily lead to misinterpretation due to the human instinct to search for regular patterns even when there may be none (Wunsch, 1999). The non-stationary RW series has the most marked trends and periodicities, whereas the decadal medians of the AR(1) series never wander far from the mean value of zero. The AR(10) and FAR(1) series are intermediate between the stationary AR(1) and the non-stationary RW cases, and show some periods containing long period variations (e.g. years 300–500). An interesting property of long-range and random walk processes is that when strong decadal trends are present, short-term variations are often less prevalent. Strong trends are often the fortuitous result of a run of changes having the same sign, and so by definition are less likely to contain rapid fluctuations.

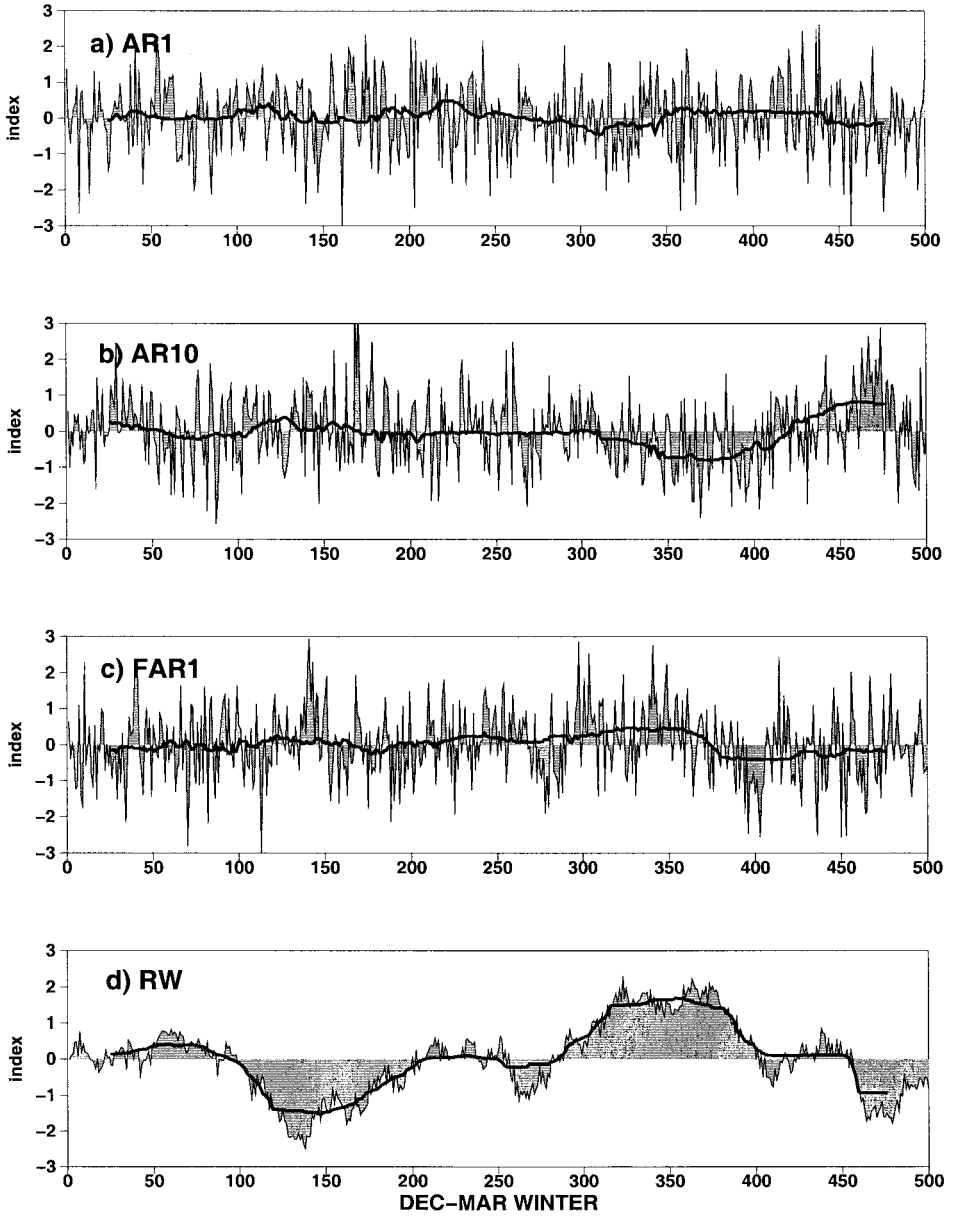

Figure 7. Some examples of 500 year long realizations simulated using the different models: (a) the red noise model AR(1); (b) the higher order model AR(10); (c) the fractionally differenced model FAR(1); and (d) the random walk model ARIMA(1, 1, 0). Long-term trends are indicated by the 50-year running medians

Similar behaviour has been noted in the NAO series (Hurrell and van Loon, 1997), and speculations have been made on possible dynamical mechanisms for such behaviour (Tourre *et al.*, 1999).

7. ASSESSMENT OF MODEL FITS TO NAO SLP INDEX

This section compares the performance of the four previously described time series models in fitting the NAO SLP index from 1864 to 1998. Fits are obtained using model parameters estimated from NAO SLP index data over the same 135-year period (Table I). In addition to the four stochastic models, statistics for two very simple benchmark forecasts are also presented: 'climatology' and 'persistence'. A 'climatology' forecast assumes that the value in the following year will be the same as the overall climatological mean, whereas a 'persistence' forecast assumes that the value the following year will be the same as in the current year (a martingale process).

7.1. Forecast verification

With the exception of the correlation study of Johansson *et al.* (1998), very few studies have been published describing forecasts of the NAO. In this section, we will focus on the skill of the four different time series models at forecasting one winter ahead the NAO SLP index from 1865 to 1998. Such verification of in-sample forecasts of past values (hindcasts) is the standard method for assessing the goodness of fit of time series models. The 1-year ahead forecast $\hat{z}_y$ for year y is compared to the observed value of z_y, and the misfit is measured by the residual $e_y = z_y - \hat{z}_y$. In-sample hindcasts differ from real-time forecasts of the unknown future in that model parameters are estimated using all the sample data, and so prior information about the values to be forecast has been included when making the forecast. In-sample hindcasts often have more forecast skill than can be realistically achieved in real application 'out-of-sample' forecasts. However, because of the shortness of the NAO SLP index series, out-of-sample forecasts could be more unreliable because of the increased uncertainty in model parameters estimated over shorter, yet independent, periods. An alternative approach is to perform cross-validated forecasts in which model parameters are estimated using all the available data except the year to be forecast (Johansson *et al.*, 1998). However, this leave-one-out approach should be used with care since it may overestimate forecast skill because of the presence of serial correlations and decadal trends in the NAO.

7.2. Goodness of fit

Figure 8 shows the model fits to the standardized NAO SLP index time series obtained by estimating the model parameters over the whole period 1864–1998. Verification statistics obtained for these fits are given in Table II. All the model fits substantially underestimate the variance of the standardized NAO SLP index—the S.D.s $\sigma(\hat{z})$ in Table II are all less than one. With the exception of the random walk model, all the fits in Figure 8 remain close to the mean value of zero. However, the RW model is more volatile than the other more stationary fits and resembles more closely a persistence forecast. None of the fits match closely the NAO SLP index, and all give residuals with S.D.s $\sigma(e)$ comparable to the S.D. of the original standardized index. For the random walk and persistence forecasts, the S.D. of the residuals actually exceeds the S.D. of the original series. The high order AR(10) model gives residuals with the smallest S.D. of 0.96, and explains only about 8% (= 100% $(1-0.96)^2$) of the total variance close to the maximum possible limit of 10% estimated by Wunsch (1999). Despite having only one and not ten parameters, the FAR(1) model gives residuals having only slightly larger S.D.s. A more robust measure of forecast error appropriate for short noisy time series is provided by the mean absolute deviation (MAD) about the median:

$$\mathrm{MAD} = \frac{1.483}{n} \sum_{y=1}^{n} |\hat{z}_y - z_y|, \tag{6}$$

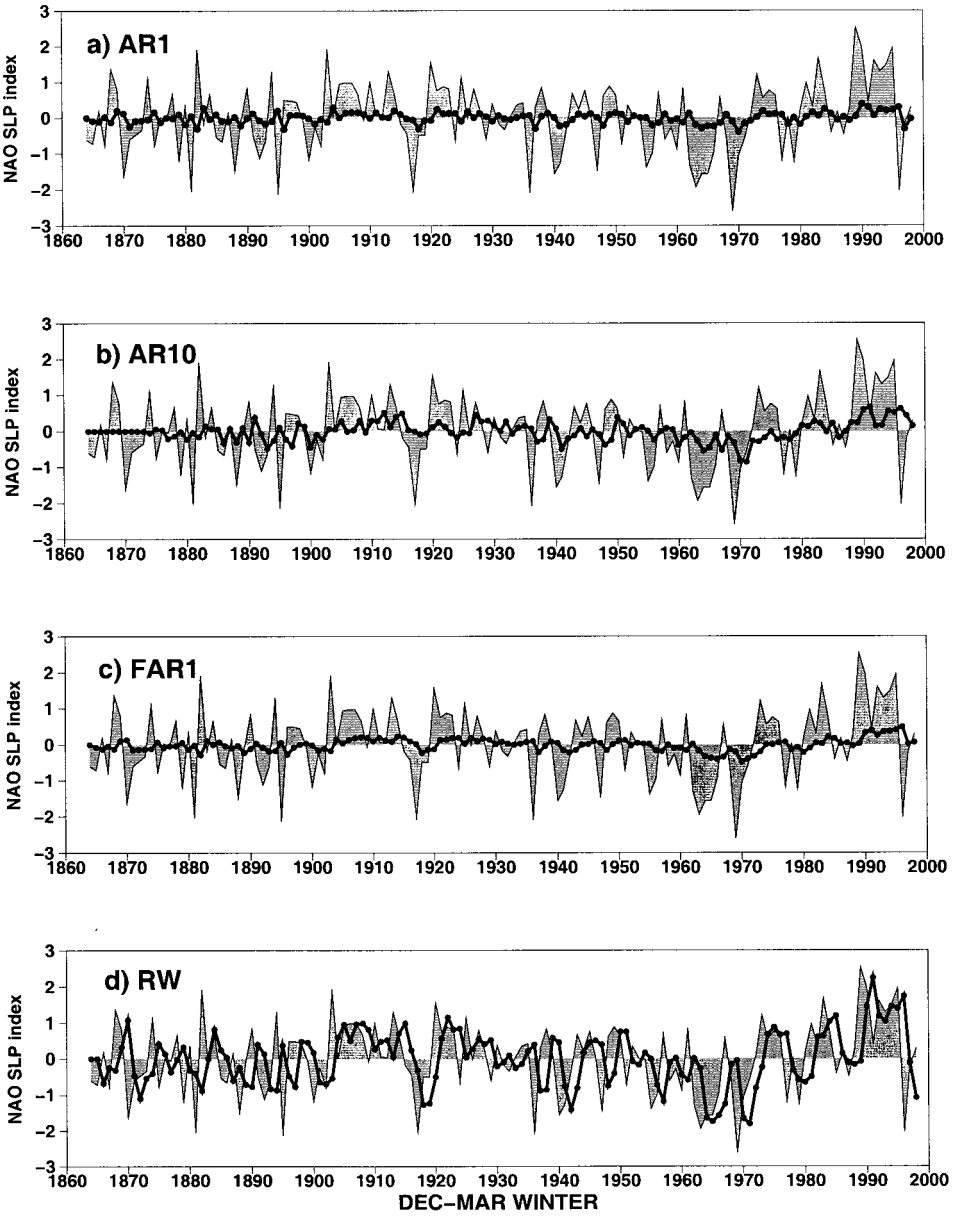

Figure 8. Model fits to the standardized NAO SLP index (in grey) obtained for: (a) the red noise AR(1) model; (b) the higher order AR(10) model; (c) the fractionally differenced FAR(1) model; and (d) the random walk ARIMA(1, 1, 0) model

which is equal to the S.D. for normally distributed residuals. From Table II, it can be seen that the random walk gives the smallest MAD of 0.93 which is much less than the S.D. of 1.15. The random walk model has a large S.D. but a small MAD because it makes risky forecasts not close to the mean and so sometimes produces much larger errors than do the other models. Because of squaring, large errors contribute significantly more to the S.D. than they do to the MAD. The FAR(1) model gives residuals with a MAD of 0.97, less than the MAD of 0.99 for the original index. Furthermore, the MAD value of 0.97 is close to the FAR(1) residual S.D. of 0.98 suggesting that the FAR(1) residuals are close to been normally distributed.

7.3. Measures of association

Table II presents two measures of association that can be used to judge the forecast skill: the product moment correlation $r(z, \hat{z})$, and the odds ratio θ. Largest correlations are obtained with the AR(10) model closely followed by the FAR(1) and RW models. All the correlations are small and are most likely predominantly due to the decadal trends in the NAO. A complementary way of assessing the skill of forecasts is to use the 'odds ratio' based on the 2×2 contingency table of the number of binary events when the index is above or below zero (Stephenson, 1999). For example, the AR(10) fit gives 39 'hits' ($z \geq 0$ and $\hat{z} \geq 0$), 23 'false alarms' ($z < 0$ and $\hat{z} \geq 0$), 29 'misses' ($z \geq 0$ and $\hat{z} < 0$), and 34 'correct rejections' ($z < 0$ and $\hat{z} < 0$). The odds ratio is equal to $39 \times 34/23 \times 29 = 1.99$ which is significantly different from unity (no skill) at 95% confidence. The odds ratio can be interpreted as the ratio of the odds of forecasting a hit to the odds of forecasting a false alarm (Stephenson, 1999). From Table II, it can be seen that the AR(10) model has the highest odds ratio followed next by the FAR(1) model. All models have odds ratios between 1 and 2 confirming that there is some slight positive association between the fits and the NAO SLP index.

7.4. Are the residuals correlated?

When a model provides a good fit to a time series, the leftover residuals are expected to be white noise having uncorrelated values. To judge this, we have calculated the aggregated variance for the residuals and used the scaling relationship $\mathrm{var}(\bar{e}_n) = \sigma^2(e) n^{\alpha}$ to estimate the exponent α. The estimate of α should be close to -1 when the residuals are uncorrelated white noise. It can be seen from Table II that the random walk and persistence models give residuals having α much less than -1 which signifies that the residuals have more power at high frequencies than at low frequencies (blue noise). This is a classic sign that the time series has been over differenced leading to anti-correlated successive values. The AR(1) and climatological forecasts give residuals having α greater than -1 which signifies that the residuals have more power at lower frequencies than at high frequencies. In other words, the residuals still contain persistence and runs that were present in the original series. The FAR(1) and AR(10) models give residuals having α close to -1 suggesting that the residuals are uncorrelated noise.

7.5. Brief summary

This section has examined various aspects of the fits provided the four different models. None of the forecasts provides a particularly close fit to the noisy NAO SLP index. The persistence and RW models give the poorest fits with correlated residuals containing more variance than that of the original index. The other models give fits which forecast 'safe' values close to the mean climatological value of zero. The AR(10) and FAR(1) models give forecasts having the largest associations with the observed index as measured by the product moment correlation and the odds ratio. Similar to what was found by Geweke

Table II. Verification statistics for model fits to the standardized NAO SLP index for the winters 1865–1998

Model	$\sigma(\hat{z})$	$\sigma(e)$	MAD(e)	$r(\hat{z}, z)$	Odds ratio	$\alpha(e)$
CLIM	0.00	1.00	0.99	–	–	−0.81
AR(1)	0.15	0.99	0.95	0.149	1.23	−0.92
AR(10)[a]	0.28	0.96	1.00	0.276	1.99	−1.00
FAR(1)	0.17	0.98	0.97	0.173	1.66	−0.99
RW	0.76	1.15	0.93	0.173	1.39	−1.72
PERS	1.00	1.30	1.28	0.149	1.23	−1.91

The one-year ahead in-sample forecasts are denoted by $\hat{z}$, and are shown in Figure 8. For good fits, the residuals $e_y = \hat{z}_y - z_y$ should be Gaussian distributed with MAD(e) close to $\sigma(e)$, and should also be independent of one another with $\alpha(e)$ close to -1. Refer to Section 7 for more discussion.
[a] The fit was performed over the slightly shorter period 1874–1998.

and Porter-Hudak (1983), the high order AR model slightly outperforms the FAR model, yet at the expense of having nine additional parameters. Therefore, FAR(1) appears to be the simplest model that provides one of the best fits to the noisy NAO SLP index.

8. CONCLUDING REMARKS

From the evidence provided by historical SLP data from 1864 to 1998, interannual variations in the wintertime NAO SLP index have a broad band spectrum which is close to being white noise. Such independence of winters is not surprising, if one believes that interannual variations in the NAO are nothing more than sampling fluctuations caused by stochastic variations in the number and intensity of North Atlantic storms each winter. However, despite being the most apparent feature, this pessimistic prospect for forecasting the NAO is not the whole story. Interestingly, there is some suggestion of 'long-range dependence' in the NAO that is required to explain the low frequency behaviour. Such behaviour is not accounted for by simple autoregressive models such as the red noise stochastic model of climate variability. Long-range fractionally integrated noise $z_y = \Delta^{-d} a_y$, having only one parameter ($d = 0.15$), provides one of the best fits to the NAO SLP index. Larger values of d such as $d = 1$ (random walk) are overdifferenced and provide less good fits. To summarize, the evidence based on the short historical record suggests that the NAO is a stationary process with long-range dependence rather than been a non-stationary random walk.

Long-range behaviour is ubiquitous and has been named 'the Joseph effect' by Mandelbrot and Wallis (1968) based on Joseph's interpretation of the Pharaoh's prophetic dream concerning climatic persistence:

> . . . there came seven years of great plenty throughout the land of Egypt. And there shall arrive after them seven years of famine. (*Genesis*, 41: 29–30).

'The Joseph effect' is used to describe time series which exhibit long-term persistence, and is not intended to signify the presence of cycles with definite periodicities. The recent series of high NAO winters in the 1990s provides a contemporary example of 'the Joseph effect'. Various mechanisms exist for generating long-range dependence such as aggregation, non-linear interactions, and self-criticality. The aggregation of an infinite number of short-range processes can lead to a process with long-range dependence and this is a plausible mechanism for the long-range dependence often seen in hydrological and financial processes (Beran, 1994). It could also be the mechanism responsible for long-range dependence in the NAO with the North Atlantic ocean integrating a large number of stochastic heat fluxes.

This study has shown that even simple stochastic models such as fractionally integrated noise are capable of generating long-range processes having significant low frequency variability and marked trends. It is not inconceivable that natural climate variability could also possess such properties simply as the result of the aggregation of many stochastic weather processes. Care should therefore be exercised when assessing and attributing causes to trends observed in climate. Simple short-range processes can underestimate the stochastic trends that may occur naturally in climate even in the absence of any deterministic influences from anthropogenic radiative forcing or dynamical changes in ocean circulation. This study has been based on the short time series of historical climate measurements that are available, and longer model generated series are currently being investigated (Stephenson and Pavan, 1999).

ACKNOWLEDGEMENTS

We thank J. Gunson, J. Hurrell, N. Kvamsto, K. Lane, T. Osborn, N. Ward, and C. Wunsch for stimulating discussions concerning this work, and J.-F. Royer and D. Besson at Météo-France for their invaluable and friendly help in retrieving historical references. DBS also wishes to thank H. Cardot, J. Beran, J.-M. Bardet, and J. Hosking for their assistance with long-range processes. We are indebted to the two reviewers for their useful remarks. This project was partially supported by the EEC Framework IV environment project STOEC.[5]

APPENDIX A. ARIMA(p, d, q) TIME SERIES MODELS

Auto-Regressive Integrated Moving Average (ARIMA) time series models form a general class of linear models that are widely used in modelling and forecasting time series (Box and Jenkins, 1976). The ARIMA(p, d, q) model of the time series $\{z_1, z_2, \dots\}$ is defined as:

$$\Phi_p(B)\Delta^d z_y = \Theta_q(B)a_y, \tag{7}$$

where B is the backward shift operator, $Bz_y = z_{y-1}$, $\Delta = 1 - B$ is the backward difference, and Φ_p and Θ_q are polynomials of order p and q, respectively. ARIMA(p, d, q) models are the product of an autoregressive AR(p) part $\Phi_p = 1 - \phi_1 B - \phi_2 B_2 - \cdots - \phi_p B^p$, an integrating part $I(d) = \Delta^{-d}$, and a moving average MA(q) part $\Theta_q = 1 - \theta_1 B - \theta_2 B^2 - \cdots - \theta_q B^q$. The parameters in Φ and Θ are chosen so that the zeros of both polynomials lie outside the unit circle in order to avoid generating unbounded processes. The difference operator takes care of 'unit root' $(1 - B)$ behaviour in the time series and for $d > 0.5$ produces non-stationary behaviour (e.g. increasing variance for longer time series). A simple example of an ARIMA model is provided by the ARIMA(1, 0, 0) first order autoregressive model $z_y = \phi_1 z_{y-1} + a_y$ that is sometimes used to model natural climate variability. All these models are 'discrete time' models suitable for modelling time series sampled at regular intervals.

NOTES

1. Freely available from http://www-sv.cict.fr/lsp/Stephen/NAO/index.html
2. Available from http://wesley.wwb.noaa.gov/reanalysis.html
3. Sometimes also referred to as the Durbin–Watson statistic.
4. Closely related modes explaining more total variance can be obtained by using either a smaller spatial domain or by using the covariance instead of the correlation matrix.
5. http://www-sv.cict.fr/lsp/Stephen/STOEC/index.html

REFERENCES

Angell, J.K. and Korshover, J. 1963. 'Harmonic analysis of the biennial zonal-wind and temperature regimes', *Mon. Weather Rev.*, **91**, 537–548.

Angell, J.K., Korshover, J. and Cotton, G.F. 1969. 'Quasi-biennial variations in the "Centres of Action"', *Mon. Weather Rev.*, **97**, 867–872.

Beran, J. 1994. *Statistics for Long-Memory Processes*, *Monographs on Statistics and Applied Probability*, *No. 61*. Chapman and Hall, 315 pp.

Bjerknes, J. 1962. 'Synoptic survey of the interaction of sea and atmosphere in the North Atlantic', *Geophys. Norvegica*, **24**(3), 115–145.

Bjerknes, J. 1964. 'Atlantic air–sea interaction', *Adv. Geophys.*, **10**, 1–82.

Bloomfield, P. 1992. 'Trends in global temperature', *Climate Change*, **21**, 1–16.

Bojariu, R. and Cotariu, R. 1996. 'Seasonal effects of Atlantic air–sea interaction on climate fluctuations over Southern and Central Europe', *Roumanian J. Meteorol.*, **3**, 1–7.

Box, G.E.P. and Jenkins, G.M. 1976. *Time Series Analysis—Forecasting and Control*. Holden-Day, California, 575 pp.

Brier, G.W. 1968. 'Long-range prediction of the zonal westerlies and some problems in data analysis', *Rev. Geophys.*, **6**, 525–551.

Clayton, H.H. 1885. 'A lately discovered meteorological cycle', *Am. Meteorol. J.*, **1**, 130, 528.

Cook, E.R., D'Arrigo, R.D. and Briffa, K.R. 1997. 'A reconstruction of the North Atlantic Oscillation using tree-ring chronologies from North America and Europe', *The Holocene*, **8**(1), 9–17.

Davis, R.E. 1976. 'Predictability of sea surface temperature and sea-level pressure anomalies over the North Pacific ocean', *J. Phys. Oceanogr.*, **6**, 249–266.

Deser, C. and Blackmon, M.L. 1993. 'Surface climate variations over the North Atlantic Ocean during winter: 1900–1989', *J. Climate*, **6**, 1743–1753.

Ebdon, R.A. 1975. 'The quasi-biennial oscillation and its association with the tropospheric circulation patterns', *Meteorol. Mag.*, **104**, 282–297.

Frankignoul, C. and Hasselmann, K. 1977. 'Stochastic climate models. II. Application to sea-surface temperature variability and thermocline variability', *Tellus*, **29**, 284–305.

Fromentin, J.-M. and Planque, B. 1996. 'Calanus and environment in the eastern North Atlantic. Part II. Influence of the North Atlantic Oscillation on *C. finmarchicus* and *C. helgolandicus*', *Mar. Ecol. Prog. Ser.*, **134**, 111–118.

Geweke, J. and Porter-Hudak, S. 1983. 'The estimation and application of long memory time series models', *J. Time Ser. Anal.*, **4**(4), 221–238.

Griffies, S.M. and Bryan, K. 1997. 'A predictability study of simulated North Atlantic multidecadal variability', *Climate Dyn.*, **13**, 459–487.

Haslett, J. and Raftery, A.E. 1989. 'Space-time modelling with long-memory dependence: assessing Ireland's wind power resource', *Appl. Stat.*, **38**(1), 1–50.
Hann, J. 1890. 'Zur Witterungsgeschichte von Nord-Grönland, Westküste', *Meteor. Zeitung*, **7**, 109–115.
Hasselmann, K. 1976. 'Stochastic climate models. Part I. Theory', *Tellus*, **28**, 473–484.
Helland-Hansen, B. and Nansen, F. 1920. 'Temperature variations in the North Atlantic ocean and in the atmosphere', *Smithsonian Misc. Collections*, **70**(4), 406.
Hildebrandsson, H.H. 1897. 'Quelques recherches sur les centres d'action de l'atmosphère', *Kungl. Svenska. Vetenskaps-Akademiens Handlingar*, **29**(3), 2–11.
Hollander, M. and Wolfe, D.A. 1973. *Nonparametric Statistical Methods*. John Wiley & Sons, New York, 528 pp.
Hosking, J.R.M. 1981. 'Fractional differencing', *Biometrika*, **68**(1), 165–176.
Hurrell, J.W. 1995. 'Decadal trends in the North Atlantic Oscillation: regional temperatures and precipitation', *Science*, **269**, 676–679.
Hurrell, J.W. 1996. 'Influence of variations in extratropical wintertime teleconnections on Northern Hemisphere temperature', *Geophys. Res. Lett.*, **23**(6), 665–668.
Hurrell, J.W. and van Loon, H. 1997. 'Decadal variations in climate associated with the North Atlantic Oscillation', *Climatic Change*, **36**, 301–326.
Hurst, H.E. 1951. 'Long-term storage capacity of reservoirs', *Trans. Am. Soc. Civil Eng.*, **116**, 770–779.
Jones, P.D., Raper, S.C.B., Bradley, R.S., Diaz, H.F., Kelly, P.M. and Wigley, T.M.L. 1986. 'Northern Hemisphere surface air temperature variations, 1851–1984', *J. Clim. Appl. Meteorol.*, **25**, 161–179.
Johansson, A., Barnston, A., Suranjana, S. and van den Dool, H. 1998. 'On the level and origin of seasonal forecast skill in Northern Europe', *J. Atmos. Sci.*, **55**, 103–127.
Kettlewell, P.S., Sothern, R.B. and Koukkari, W.L. 1999. 'U.K. wheat quality and economic value are dependent on the North Atlantic Oscillation', *J. Cereal Sci.*, in press.
Kushnir, Y. 1994. 'Interdecadal variations in North Atlantic sea surface temperature and associated atmospheric conditions', *J. Climate*, **7**, 142–157.
Lamb, H.H. 1972. *Fundamentals and Climate Now*, *vol I*, *Climate: Present*, *Past and Future*. Methuen, 613 pp.
Landsberg, H.E., Mitchell, J.M., Crutcher, H.L. and Quinlan, F.T. 1963. 'Surface signs of the biennial atmospheric pulse', *Mon. Weather Rev.*, **91**, 10–12, 549–556.
Loewe, F. 1937. 'A period of warm winters in western Greenland and the temperature see-saw between western Greenland and Europe', *Q. J. R. Meteorol. Soc.*, **63**, 365–371.
Madden, R.A. 1976. 'Estimates of the natural variability of time-averaged sea-level pressure', *Mon. Weather Rev.*, **104**, 942–952.
Mandelbrot, B.B. and Wallis, J.R. 1968. 'Noah, Joseph, and operational hydrology', *Water Resources Res.*, **4**, 909–918.
Meinardus, W. 1900. 'Einige Beziehungen zwischen der Witterung und den Ernte-Erträgen in Nord-Deutschland', *Verhadl. des VII. Int. Geographen Kongresses in Berlin (1899).*
Meehl, G.A. 1997. 'The South Asian Monsoon and the Tropospheric Biennial Oscillation (TBO)', *J. Climate*, **10**, 1921–1943.
Montanari, A., Rosso, R. and Taqqu, M.S. 1996. 'Some long-run properties of rainfall records in Italy', *J. Geophys. Res.*, **101**(D23), 29, 431–429, 438.
Rogers, J.C. 1984. 'The association between the North Atlantic Oscillation and the Southern Oscillation in the Northern Hemisphere', *Mon. Weather Rev.*, **112**, 1999–2015.
Rogers, J.C. 1990. 'Patterns of low-frequency monthly sea-level pressure variability (1899–1986) and associated wave cyclone frequencies', *J. Climate*, **3**, 1364–1379.
Saravanan, R. and McWilliams, J.C. 1997. 'Stochasticity and spatial resonance in interdecadal climate fluctuations', *J. Climate*, **10**, 2299–2320.
Stephenson, D.B. 1999. 'Use of the odds ratio for diagnosing forecast skill', *Weather Forecast.* (in press).
Stephenson, D.B. and Pavan, V. 1999. 'An exploration of NAO-related temperature variations in 17 global coupled ocean atmosphere models', *Climate Dyn.* (to be submitted).
Teisserenc de Bort, M.M. 1883. 'Étude sur l'hiver de 1879–80 et recherches sur l'influence de la position des grands centres d'action de l'atmosphère dans les hivers anormaux', *Ann. Soc. Météor. France*, **31**, 70–79.
Thompson, D.W.J. and Wallace, J.M. 1998. 'Observed linkages between Eurasian surface air temperatures, the North Atlantic Oscillation, Arctic sea-level pressure and the stratospheric polar vortex', *Geo. Res. Letts.*, **25**, 1297–1303.
Tourre, Y.M., Rajagopalan, B. and Kushnir, Y. 1999. 'Dominant patterns of climate variability in the Atlantic Ocean region during the last 136 years', *J. Climate* (submitted).
Trenberth, K.E. 1980. 'Atmospheric quasi-biennial oscillations', *Mon. Weather Rev.*, **108**, 1370–1377.
Trenberth, K.E. and Shin, W.-T. K. 1984. 'Quasi-biennial fluctuations in sealevel pressures over the Northern Hemisphere', *Mon. Weather Rev.*, **112**, 761–777.
Trenberth, K.E. 1990. 'Recent observed interdecadal climate changes in the Northern Hemisphere', *Bull. Am. Meteorol. Soc.*, **71**(7), 988–993.
van Loon, H. and Rogers, J.C. 1978. 'The seesaw in winter temperatures between Greenland and Northern Europe. Part I: general description', *Mon. Weather Rev.*, **106**, 296–310.
von Neumann, J. 1941. 'Distribution of the ratio of the mean square successive difference to the variance', *Ann. Math. Stat.*, **12**, 367–395.
von Neumann, J. 1942. 'A further remark concerning the distribution of the ratio of the mean square successive difference to the variance', *Ann. Math. Stat.*, **13**, 86–88.
Walker, G.T. 1924. 'Correlations in seasonal variations of weather IX', *Mem. Indian Meteorol. Dept.*, **24**, 275–332.
Wallace, J.M., Zhang, Y. and Bajuk, L. 1996. 'Interpretation of interdecadal trends in Northern Hemisphere surface air temperature', *J. Climate*, **9**, 249–259.
Wunsch, C. 1999. 'The interpretation of short climate records', *Bull. Am. Meteorol. Soc.*, **80**(2), 245–255.

