University of Reading

A regime view of the North Atlantic Oscillation and its response to anthropogenic forcing

Article

Published Version

Woollings, T., Hannachi, A., Hoskins, B. and Turner, A. ORCID: https://orcid.org/0000-0002-0642-6876 (2010) A regime view of the North Atlantic Oscillation and its response to anthropogenic forcing. Journal of Climate, 23 (6). pp. 1291-1307. ISSN 1520-0442 doi: 10.1175/2009JCLI3087.1 Available at https://centaur.reading.ac.uk/8174/

It is advisable to refer to the publisher's version if you intend to cite from the work. See Guidance on citing.
Published version at: http://dx.doi.org/10.1175/2009JCLI3087.1
To link to this article DOI: http://dx.doi.org/10.1175/2009JCLI3087.1

Publisher: American Meteorological Society

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the End User Agreement.

www.reading.ac.uk/centaur

University of Reading

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

A Regime View of the North Atlantic Oscillation and Its Response to Anthropogenic Forcing

TIM WOOLLINGS AND ABDEL HANNACHI

Department of Meteorology, University of Reading, Reading, United Kingdom

BRIAN HOSKINS

Department of Meteorology, University of Reading, Reading, and Grantham Institute, Imperial College, London, United Kingdom

ANDREW TURNER

NCAS–Climate, and Department of Meteorology, University of Reading, Reading, United Kingdom

(Manuscript received 10 February 2009, in final form 15 September 2009)

ABSTRACT

The distribution of the daily wintertime North Atlantic Oscillation (NAO) index in the 40-yr ECMWF Re-Analysis (ERA-40) is significantly negatively skewed. Dynamical and statistical analyses both suggest that this skewness reflects the presence of two distinct regimes—referred to as ''Greenland blocking'' and ''subpolar jet.'' Changes in both the relative occurrence and in the structure of the regimes are shown to contribute to the long-term NAO trend over the ERA-40 period.

This is contrasted with the simulation of the NAO in 100-yr control and doubled CO_2 integrations of the third climate configuration of the Met Office Unified Model (HadCM3). The model has clear deficiencies in its simulation of the NAO in the control run, so its predictions of future behavior must be treated with caution. However, the subpolar jet regime does become more dominant under anthropogenic forcing and, while this change is small it is clearly statistically significant and does represent a real change in the nature of NAO variability in the model.

1. Introduction

The North Atlantic Oscillation (NAO) is the dominant pattern of atmospheric variability in the extratropical Northern Hemisphere winter and was one of the earliest patterns of variability to be discovered [see, e.g., Wanner et al. (2001) and Stephenson et al. (2004) for a historical perspective]. The earliest proposed mechanisms saw the NAO as a coupled mode of climate variability between the North Atlantic surface ocean and the overlying atmosphere (Bjerknes 1964). However, a more recent view is that it is essentially an internal mode of atmospheric variability (Hurrell et al. 2002), which has a characteristic time scale of around 10 days (Feldstein 2000) but does exhibit long-range dependence on the interannual time scale (Stephenson et al. 2000).

While the NAO is often represented by its associated pattern of pressure or geopotential height anomalies, it can be viewed as essentially the signal of combined variations in the strength and orientation of the Atlantic storm track and the associated eddy-driven jet (Thompson et al. 2002; Vallis and Gerber 2008). To emphasize this interpretation of the NAO, we begin not with the familiar anomaly pattern, but with an illustration of the underlying changes in the jet stream. Figure 1 shows the upper-level winds for positive and negative NAO days. During positive days there is a strong subpolar jet oriented southwest–northeast toward Europe. This is termed the eddy-driven jet as it owes its existence to the mean flow forcing of transient eddies. As described by Ambaum et al. (2001), it is clearly separated from the subtropical jet that is developing over the subtropical North Atlantic. On negative NAO days, in contrast, the two jet

Corresponding author address: Tim Woollings, Dept. of Meteorology, University of Reading, Earley Gate, P.O. Box 243, Reading, RG6 6BB, United Kingdom.
E-mail: t.j.woollings@reading.ac.uk

DOI: 10.1175/2009JCLI3087.1

© 2010 American Meteorological Society

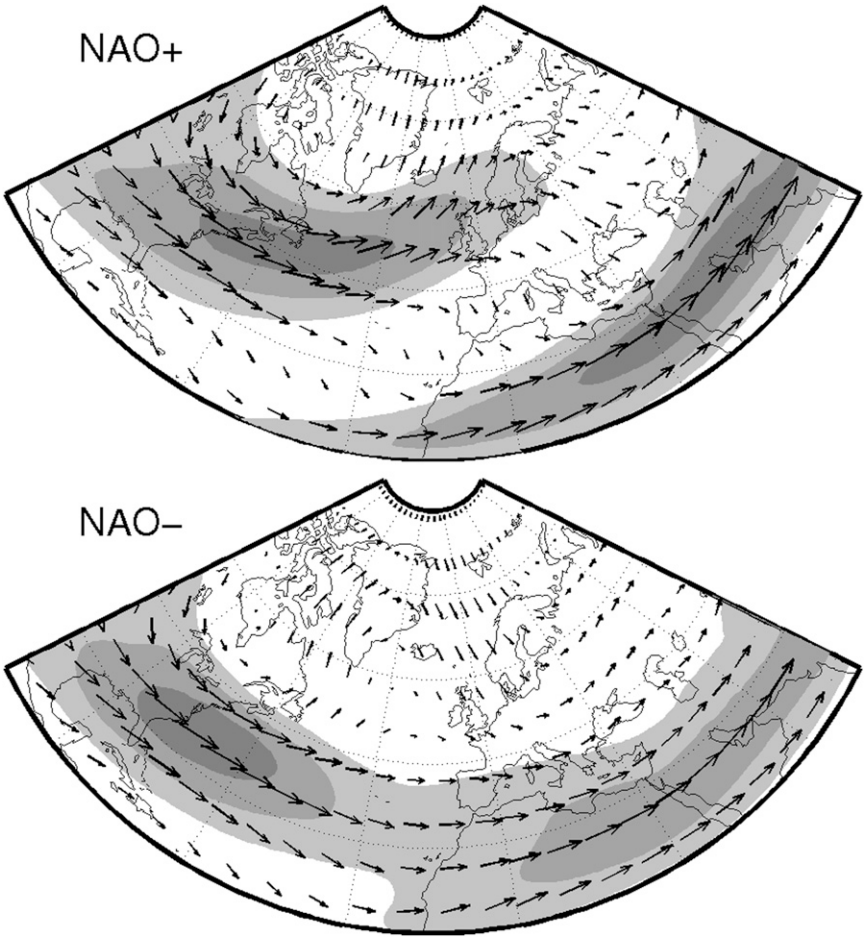

FIG. 1. Composites of the winter (DJF) 300-hPa wind field during positive and negative NAO days, using a threshold of one standard deviation of the standard NAO index defined in section 3. Isotachs are shaded at 20, 30, and 40 m s^{-1}.

streams have merged to form one broad, continuous jet across the Atlantic.

Several studies have suggested that these jet stream variations arise as a result of mean flow forcing associated with the breaking of transient, synoptic-scale Rossby waves (e.g., Benedict et al. 2004; Franzke et al. 2004; Strong and Magnusdottir 2008). Wave breaking on the equatorward side of the jet tends to be anticyclonic, following the ambient background shear. This leads to poleward eddy fluxes of zonal momentum that act to push the jet to the north. Similarly, cyclonic wave breaking dominates on the poleward side of the jet, and the associated momentum fluxes push it to the south. This suggests a three-state view of the NAO, comprising a background state and positive and negative NAO regimes. In contrast, Woollings et al. (2008, hereafter W08) suggested that there may only be two distinct regimes, corresponding to blocked and zonal flows. In this view the mean state is simply some weighted average of the two regimes. Variations in the NAO can arise from variations within each regime, such as changes in the strength of the zonal flow when there is no blocking or in changes of the residence frequency of the regimes, that is, the frequency of blocking occurrence. Some evidence of the asymmetry of the NAO has already been given by Cassou et al. (2004) and Blessing et al. (2005).

There is much interest in the suggestion that atmospheric variability exhibits preferred flow states, or regimes (Kimoto and Ghil 1993; Cheng and Wallace 1993; Palmer 1999; Christiansen 2005a; Hannachi 2007). This interest is partly driven by the potential for regime behavior to enhance the understanding and predictability of atmospheric variability on a time scale of weeks or even longer (Straus et al. 2007). The issue of climate change has further enhanced interest in the subject, as it is suggested that the response to anthropogenic forcing may be felt as a change in the residence frequency of the most dominant regimes (e.g., Palmer 1999; Corti et al.

1999; Hsu and Zwiers 2001; Terray et al. 2004; Keeley et al. 2008). Preferred flow regimes have been identified from idealized low-order models (Charney and Devore 1979) and intermediate complexity models (Egger 1981; Legras and Ghil 1985; Crommelin 2004; Franzke et al. 2008). The capability of complex climate models to capture preferred large-scale flow patterns, in simulations without anthropogenic forcing, has been demonstrated by Haines and Hannachi (1995), Hannachi (1997), Branstator and Berner (2005), and Berner (2005) and, in simulations with anthropogenic forcing, by Hsu and Zwiers (2001), Monahan et al. (2000), and Hannachi and Turner (2008). Flow regimes from reanalyses have also been diagnosed by many authors, see, for example, Mo and Ghil (1987), Monahan et al. (2001), Crommelin (2004), and Hannachi (2007, hereafter H07) and references therein, but the issue is still under discussion, given the possibility of the presence of multiplicative noise (Sura et al. 2005) and the relatively small sample size, particularly for monthly and longer time scales (Wallace et al. 1991; Stephenson et al. 2004); see H07 for more discussion. The tools used to identify preferred regimes range from dynamical, based on quasi-stationarity of the preferred flow pattern, to statistical, based on clustering or modeling the probability density function (PDF) of the system within its state space. See H07 and Handorf et al. (2009) for details and more references.

There has been much work on regime behavior using statistical methods and on the wave-breaking theories using dynamical methods, but there have been relatively few attempts to combine these two approaches. Here we present a first attempt at such an analysis. We show that the daily NAO index is negatively skewed and, as in H07, we interpret skewness in terms of a mixture model approach used to model the NAO PDF. Both the mixture model and the dynamical wave-breaking index of W08 suggest that this skewness could be explained by the presence of two distinct preferred flow regimes. We examine the recent trend in the NAO from this regime perspective and then progress to examine data from simulations of a coupled general circulation model (GCM). We assess the ability of the model to represent this regime behavior and also investigate whether the model's response to anthropogenic forcing involves a modification of the regime structure.

2. Data and methods

a. Reanalysis data

We use sea level pressure (SLP) and 500-hPa geopotential height (Z_{500}) data from 44 complete winters [December–February (DJF) 1957/58–2000/01] from the 40-yr European Centre for Medium-Range Weather Forecasts Re-Analysis (ERA-40) (Uppala et al. 2005). The data is defined on a 2.25° × 2.25° grid covering the region north of 20°N. Monthly anomalies are calculated by subtracting the mean of the respective calendar month, and daily anomalies are calculated by removing a smoothed seasonal cycle, which is derived by averaging the daily values over all years and then smoothing with a discrete cosine transform, retaining only the mean and the lowest two Fourier frequencies.

b. GCM data

The model data are derived from simulations of a 30-layer version of the third climate configuration of the Met Office Unified Model (HadCM3). This is a coupled atmosphere–ocean model with an atmosphere resolution of 2.5° × 3.75°. We use output from two 100-yr simulations: a preindustrial control and an equilibrium simulation with doubled atmospheric carbon dioxide (CO_2). These model runs were analyzed in Hannachi and Turner (2008), where more details on the model and the experimental configuration can be found. Monthly SLP and Z_{500} anomalies are calculated by subtracting the mean of the respective calendar month, as above. Daily anomalies will be used to characterize the change in the distribution of the NAO index, so it is desirable to maintain the mean difference between the two runs. To achieve this a smoothed seasonal cycle was derived for the control run, as for ERA-40, and the anomalies for both experiments were calculated by subtracting this from the data. This approach is justified since there is no significant change in seasonality between the two experiments (Hannachi and Turner 2008). All data were analyzed on the model grid using all points north of 20°N. Unfortunately, the data needed to calculate the wave-breaking index was not available for the runs, so only the mixture model can be applied here. Some other blocking indices require less specialized data but have not been formulated to identify the high-latitude events that we focus on here.

c. The wave-breaking index

W08 identified persistent blocking-like events over Greenland, which they referred to as Greenland blocking episodes (GBEs). These arise from cyclonic wave-breaking events near the start of the Atlantic storm track, and correspond to the cyclonic breaking events described by Benedict et al. (2004) and others. The GBEs comprise one of the two NAO regimes, suggested by W08, with the other regime representing more zonal flow. Here we use W08's decomposition of all winter days into two subsets: the set of all GBE days and the set of all other (non-GBE) days.

The GBEs were originally identified using the two-dimensional index described by Berrisford et al. (2007). This index identifies blocking episodes via the associated wave breaking, by searching for a reversal in the meridional contrast in potential temperature θ on the dynamical tropopause (θ_{PV2}; the 2-PV-unit surface). At each point, θ_{PV2} is averaged over two boxes of 5° longitude × 15° latitude to the north and south of the point. When the value of the northern box minus the southern box becomes larger than zero, a reversal is defined. Temporal and spatial constraints are then applied to ensure that the events identified are large scale, quasi-stationary, and persistent (lasting at least 5 days); these are then termed episodes. See Berrisford et al. (2007) or W08 for more details. This is referred to in general as a wave-breaking index, and identifies events in mid-latitudes classed as blocking, and events on the poleward side of the storm tracks generally termed high-latitude blocking.

W08 found that high-latitude blocking in the region 50°–60°N, 30°–70°W was particularly clearly related to the NAO. A GBE is said to occur whenever a wave-breaking episode is seen anywhere within this region.[1] Greenland blocking is identified on 1608 out of a total of 3960 DJF days, so the non-GBE regime is said to occur on 2352 days.

d. The mixture model

We use the mixture model of H07 to estimate the probability density function of the NAO. This method relies on a general result, which states that any probability density function $f(\mathbf{x})$ can be decomposed as closely as desired by a weighted average, or a mixture of multivariate Gaussian density functions (Anderson and Moore 1979); that is,

$$f(\mathbf{x}) = \sum_{k=1}^{c} \alpha_k g_k(\mathbf{x}, \boldsymbol{\Sigma}_k, \boldsymbol{\mu}_k), \tag{1}$$

where $\alpha_1, \ldots, \alpha_c$ are the c mixing proportions of the mixture model and satisfy

$$0 < \alpha_k < 1, \quad \text{for} \quad k = 1, \ldots c; \quad \sum_{k=1}^{c} \alpha_k = 1, \tag{2}$$

and $\boldsymbol{\mu}_k$ and $\boldsymbol{\Sigma}_k$ are, respectively, the mean and the covariance matrix of the kth, $k = 1, \ldots c$, multivariate normal density function, g_k:

$$g_k(\mathbf{x}, \boldsymbol{\Sigma}_k, \boldsymbol{\mu}_k) = (2\pi)^{-d/2} |\boldsymbol{\Sigma}_k|^{-1/2} \exp\left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_k)^{\mathrm{T}} \boldsymbol{\Sigma}_k^{-1} (\mathbf{x} - \boldsymbol{\mu}_k)\right], \tag{3}$$

where d is the state space dimension; see also Hannachi and O'Neill (2001) and H07 for more details. The $[c(d + 1)(d + 2) - 2]/2$ unknown parameters of model (1) are obtained using the expectation–maximization (EM) algorithm (Everitt and Hand 1981; McLachlan and Basford 1988; Hannachi and O'Neill 2001). The number of significant components of this model is estimated using arguments based on order statistics. The model starts by fitting a two-component model to the data, then repeatedly adds an extra component until the latest component does not pass the significance test. The method can be applied directly to daily data without the need to estimate an independent sample size (see H07 for more details). The mixture model is particularly well suited to this application since the information on the wave-breaking index is available for comparison. The model does not definitively allocate a particular regime to any given day but, as it models the PDF of the NAO, it can be used to give the probability that the day is contained in each regime, as is done in section 3d.

3. NAO analysis using ERA-40

a. Skewness of the NAO

The principal pattern of the NAO used here was defined by the first empirical othogonal function (EOF) of the ERA-40 monthly mean Z_{500} anomalies over an Atlantic sector (20°–90°N, 90°W–90°E). The data are weighted by the square root of the cosine of the latitude prior to calculating the EOFs, as is conventional. The pattern is shown in Fig. 2 and features the two familiar Atlantic centers of action and only weak values elsewhere. This pattern is associated with 26% of the variance in the monthly mean field over this sector. In the rest of the paper we focus attention on the NAO as defined in this manner, but in this subsection we also show results using other NAO indices to demonstrate robustness to the choice of definition.

A daily NAO index is defined by the area-weighted projection of the daily anomalies onto the NAO pattern. The index is then normalized by its standard deviation. Although the mixture model provides a model of the NAO PDF, we also use a kernel method with the normal parameter, $1.06\,\sigma n^{-1/5}$ (Silverman 1981), to obtain a general estimate of the NAO PDF. Here σ is the standard deviation of the time series and n is the sample size. Figure 3a shows the kernel estimate of the NAO PDF,

[1] Note that the latitude associated with an event is the latitude at which the meridional gradient reverses, rather than the latitude of the anticyclone.

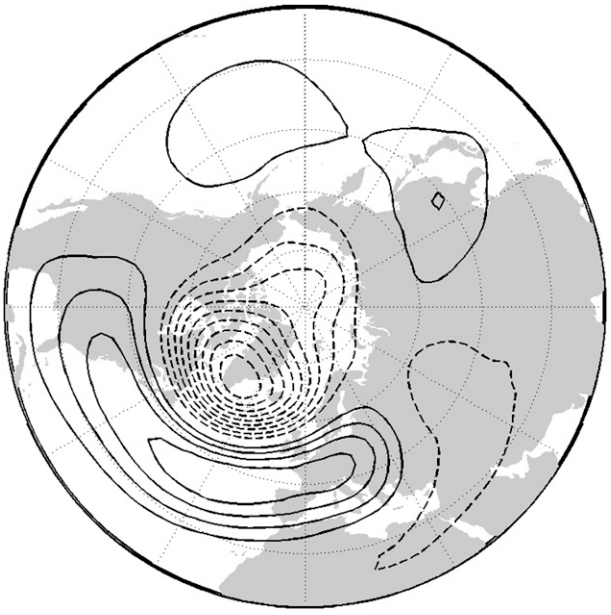

FIG. 2. The winter NAO pattern, defined as the first EOF of monthly mean Z_{500} over 20°–90°N, 90°W–90°E. The pattern is shown by regressing the monthly anomalies onto the principal component time series, with a contour interval of 10 m per standard deviation. Negative contours are dashed and the zero contour is omitted.

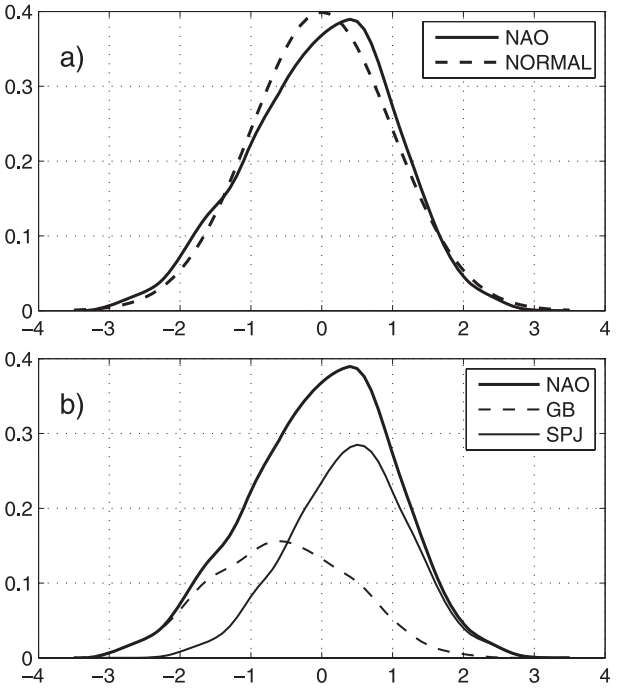

FIG. 3. (a) PDF of the daily NAO index compared to a normal distribution. (b) The same PDF split into the GBE and non-GBE days of W08.

which has a clear negative skew. The skewness is −0.23, which can be compared to the standard error on skewness given by $\sqrt{6/N}$, where N is an estimate of the number of independent samples in the data. Feldstein (2000) calculated the e-folding time scale of the NAO as 9.5 days and assuming an independent sample every 9.5 days gives a standard error of 0.12, so the observed skewness is almost twice the standard error (in 95% of similar samples from a normal distribution the absolute value of the skewness will be less than two standard errors). If a 10-day low-pass Lanczos filter is applied to the daily fields before calculation of the NAO index, the skewness intensifies to −0.27, but we prefer to use unfiltered data so that no time scale is imposed. The skewness is fairly robust with respect to the method used to derive the pattern. For example, if the NAO is defined as the first rotated EOF of monthly mean Z_{500}, following Hannachi et al. (2007), the skewness of the daily index is −0.22. Using SLP instead of Z_{500} gives a lower skewness of −0.12 for both the Atlantic and rotated EOF methods. If the NAO pattern used by the NOAA Climate Prediction Center[2] (NOAA/CPC) is used instead, the skewness of the daily index is −0.27. Finally, a qualitatively similar skewness is also evident in the NAO distribution derived from the National Centers for Environmental Prediction–National Center for Atmospheric Research reanalysis shown in Coppola et al. (2005, Fig. 2c).

The NAO index is often defined by the difference in SLP between observing stations in Iceland and the Azores (or nearby). To examine the skewness of such an index, a pseudo station index was derived from the daily ERA-40 anomalies using grid points at (65°N, 20°W) and (38°N, 26°W). To filter out very high frequency variations the SLP anomalies were averaged over consecutive, nonoverlapping 5-day periods. Following Hurrell (1995), the time series at the two points were then normalized prior to differencing, though this made no appreciable difference to the skewness. The index was then normalized and its distribution is shown in Fig. 4. The distribution exhibits a strong skewness of −0.39 with a particularly well-defined region of enhanced probability in the negative NAO phase that is suggestive of regime behavior. If the same approach is applied to Z_{500} data, the resulting distribution has a similar skewness of −0.41.

To examine the distribution of the NAO over a longer period than that of the reanalysis it is necessary to use monthly reconstructed data. Monthly data are not ideally

[2] CPC calculate a monthly NAO index from rotated principal component analysis, which is available online (http://www.cpc.noaa.gov/data/teledoc/telecontents.shtml). The monthly Z_{500} anomalies were regressed onto this index to derive the NAO pattern, which was then used to derive a daily NAO index as before.

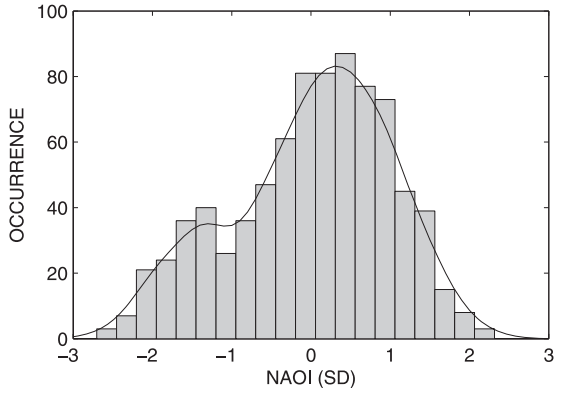

FIG. 4. Distribution of the pentad psuedo station NAO index described in section 3, shown both as a histogram and using the kernel estimation.

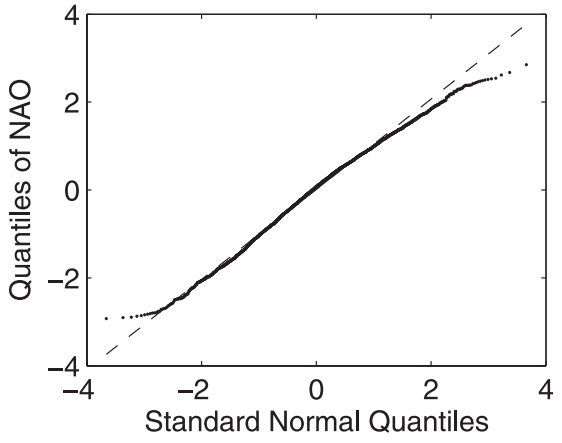

FIG. 5. Quantile–quantile (QQ) plot of the daily NAO distribution against a standard normal. The quantiles of the NAO distribution are simply plotted against the corresponding quantiles of a normal distribution.

suited for examination of non-Gaussian or regime behavior for several reasons. In addition to reducing the sample size of observational data, the flow variations in question evolve on shorter time scales, so the temporal averaging brings the data closer to normality (Teng et al. 2004). Jones et al. (1997) developed a monthly NAO index from station data, which is available online (http://www.cru.uea.ac.uk/cru/data/nao.htm) for the years 1821–2000. The distribution of this monthly NAO index over the winter (DJF) months has a skewness of −0.22 (compared to a standard error of 0.11), so in fact even the monthly index is significantly skewed.

As a final test of the sensitivity of the skewness to the choice of index, we also used a wind-based index of the NAO. Given the interpretation, outlined in the introduction, that the NAO is essentially the signal of variations in the latitude and strength of the Atlantic eddy-driven jet stream, it seems desirable to identify the NAO directly from the wind, rather than pressure or geopotential height fields. Geopotential height differs systematically from the streamfunction, so it is important to check that the skewness does not arise because of this. The NAO emerges as the first rotated EOF of the monthly mean Northern Hemisphere 500-hPa zonal wind from ERA-40. A daily index is derived, as before, by projecting the rotated EOF pattern onto the daily anomalies. The resulting distribution has a skewness of −0.32, which is much larger than the standard error of 0.12.

The departures from Gaussian behavior of the tails of the NAO distribution are hard to distinguish in Fig. 3a, so we show a quantile–quantile (QQ) plot of our standard daily NAO index versus a normal distribution in Fig. 5. The departure of the right-hand end of the curve below the diagonal shows that the positive tail of the NAO distribution is thinner than that of the normal distribution. Similarly, the departure of the left-hand end of the curve above the diagonal also reflects a thinner tail in the NAO distribution than in the normal distribution. Between around −1.5 and −3 standard deviations, the NAO PDF is above the normal, but beyond around −3 standard deviations the normal tail becomes thicker. This feature is very hard to see in Fig. 3a but is clear in Fig. 5. This shows that the negative skewness of the NAO is a feature of the moderate values of the distribution rather than the tail.

The importance of the moderate values is also verified by using the Yule–Kendall index of skewness:

$$\gamma_{\mathrm{YK}} = \frac{(q_{75} - q_{50}) - (q_{50} - q_{25})}{q_{75} - q_{25}}, \tag{4}$$

where q_i is the ith percentile. This measure of skewness is resistant to small changes in the extreme values of the distribution. For our standard daily NAO index, $\gamma_{\mathrm{YK}} = -0.086$. To test the significance, we compared this to the values of γ_{YK} seen in time series of the same length generated by a first-order autoregressive (AR1) model designed to model the NAO index (see section 3c). In a two-sided test with a sample of 10 000 trials only 0.03% of the trials exhibited γ_{YK} values as large as this (and 0.3% of the trials had conventional skewness values at least as large as that observed).

To summarize, several different NAO indices have been defined using different datasets, time scales, methods, and flow variables and all exhibit negative skewness to some extent, often with high significance. We now focus on the Z_{500} Atlantic EOF index, which gives a somewhat midrange skewness, and suggest that this skewness could be a signature of regime behavior.

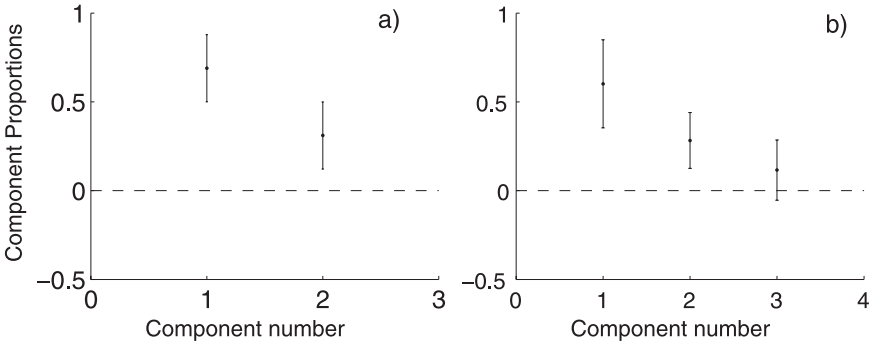

FIG. 6. The component proportions α when models of two and three regimes are fitted to the daily NAO index, along with their 95% confidence intervals.

b. A regime explanation for the skewness

W08 suggested that there are only two distinct regimes of the NAO, comprising the GBE and non-GBE days. To illustrate this the daily Z_{500} (Atlantic EOF) NAO distribution is shown in Fig. 3b, decomposed into the two subsets as described in section 2c. These two distributions exhibit less skewness: the GBE distribution has a skew of just 0.1 compared to a standard error of 0.19, while the non-GBE distribution has a skew of −0.17, similar to the standard error, which is 0.16 in this case. The skew in the NAO index could therefore be explained by the existence of two distinct regimes, both of which have distributions that are not significantly different from Gaussians.

This regime decomposition is now compared to that identified by the mixture model when applied to the same NAO distribution. Figure 6 shows the component proportions, α, when Gaussian mixture models of two (panel a) and three (panel b) components are fitted. The two-component model (Fig. 6a) identifies two regimes that are both significant, showing that at least two components can be fitted to the NAO data. To check this, Fig. 6b shows the same plot but with three components instead. It is clear that only two components should be fitted to the data as the third component is not significant. Since each individual Gaussian component is interpreted as a regime, the mixture model indicates that the full NAO distribution comprises only two distinct regimes. The distributions of these regimes are shown in Fig. 7 and are very similar to those identified by the wave-breaking index, shown in Fig. 3b. The mixture model has also been applied to the other daily NAO indices described above with similar results. Table 1 summarizes the results of the mixture model applied to the EOF-based NAO indices used here, both in Z_{500} and SLP.

Christiansen (2007) issued a useful caution on the use of a mixture model to identify regimes in atmospheric flow. However, there is good agreement between the results of the mixture model and the dynamical wave-breaking index. This adds physical insight and increases confidence compared to the results of a statistical analysis alone. Christiansen (2007) also gave an example in which the number of regimes identified by a mixture model increased as the sample size was increased. However, the method used here has been designed precisely to avoid this problem (see H07). We have also tested this specifically by applying the mixture model to 1000-, 2000-, 3000-, and 4000-day samples of the NAO index, and the results are very similar in each case (not shown).

The results of both the wave-breaking index and the mixture model support the suggestion of W08 that there are just two distinct regimes of NAO variability. We will refer to these as Greenland blocking and the subpolar jet. The mixture model results in a more compact decomposition with less overlap between the regimes. The population of the Greenland blocking regime comprises 31% of all days according to the mixture model but 41%

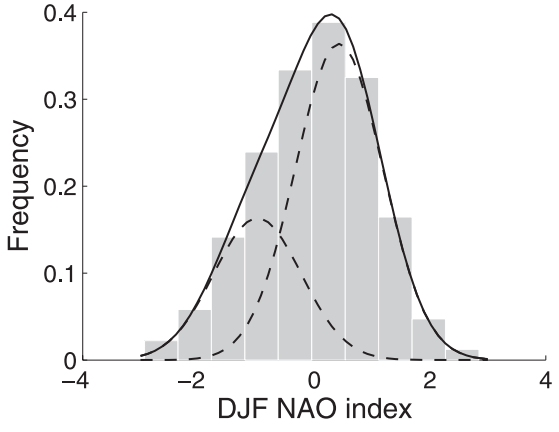

FIG. 7. Distribution of the daily NAO index along with the PDFs of the two-component mixture model fit (solid line) and of the individual components (dashed lines).

TABLE 1. Number of Gaussian components in the mixture model for different significance levels and different NAO indices from ERA-40 for the whole dataset and the first and second halves of the dataset. The indices are SLP1, Z1— EOF1 over the Atlantic sector for SLP and Z_{500}, respectively; and SLP2, Z2—rotated EOF1 for SLP and Z_{500}, respectively.

	SLP1			SLP2			Z1			Z2		
Significance (%)	All	First	Second	All	First	Second	All	First	Second	All	First	Second
5	2	2	1	2	2	2	2	2	2	2	2	2
2.5	1	2	1	2	2	1	2	2	1	2	2	2
1	1	1	1	1	2	1	2	2	1	2	2	1

according to the wave-breaking index, suggesting that the wave-breaking definition may slightly overestimate the occurrence of blocking. Figure 3 shows that the distribution of the Greenland blocking regime is very wide, containing several positive NAO days. A composite of these days features a cyclonic anomaly off the coast of Newfoundland and an anticyclonic anomaly over Scandinavia (not shown). This suggests that on these days the gradient reversal lies between these two features, that is, in the mid-Atlantic at the far eastern end of the region used to identify Greenland blocking.

The results of the mixture model have been used to derive maps of the two regimes for comparison with those in W08. These maps were derived in a similar way to a statistical expectation: by integrating the Z_{500} anomalies for all days, weighted by the probability that each day is in a given regime. The probabilities are estimated using the daily NAO index and the component PDFs in Fig. 7. The resulting maps are shown in Fig. 8 and these are very similar to versions generated by other methods, such as compositing days close to the centers of the two regimes. By construction the anomaly patterns have the same structure but opposing sign, and the magnitude of the anomalies is in line with the distance of each regime center to the origin. There is an interesting northwest–southeast tilt to the anomaly patterns. The difference in the full flow between the two regimes is especially clear just west of the British Isles, with a strong gradient in geopotential in the subpolar jet regime and a ridge in the Greenland blocking regime.

The existence of a distinct regime in negative NAO space is in agreement with the result obtained from the nonlinear principal component analysis of Monahan et al. (2000), which identified a distinct episodic negative NAO regime [though see Christiansen (2005b) for a critique of this method]. Regimes similar to Greenland blocking have been identified in regime analyses by Cheng and Wallace (1993), Vautard (1990), Kimoto and Ghil (1993), Smyth et al. (1999), and others. The importance of Greenland blocking for the NAO was also demonstrated by Croci-Maspoli et al. (2007), who showed that when all blocking days are removed from ERA-40, the NAO is no longer the first EOF over the Atlantic/European region. The decomposition of the NAO into Greenland blocking and subpolar jet regimes is also in agreement with the view of Luo et al. (2007).

Greenland blocking is, in fact, visible in the composite of the full wind field for the negative NAO days in Fig. 1. The jet stream proceeds more or less zonally across the Atlantic, as described before, but there is also a clear anticyclonic flow over the northern North Atlantic and Greenland, and this is the signature of the blocking. There is clearly an actual anticyclone there in the full field, not just an anticyclonic anomaly. It is the distinction between blocking and transient wave breaking that explains this anticyclonic flow. If the cyclonic wave breaking was transient, then the flow here would presumably be cyclonic on average. The reason for the anticyclonic flow is that large masses of relatively low potential vorticity

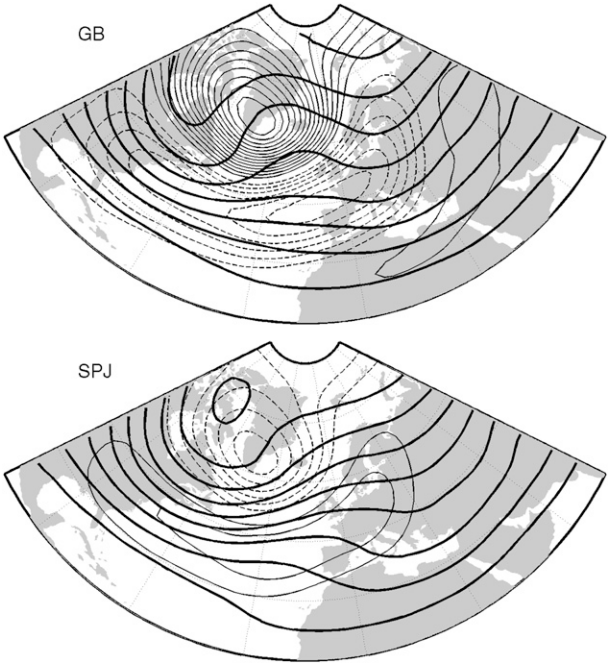

FIG. 8. Maps of the expected value of Z_{500} for the two regimes identified by the mixture model. Thick contours show the full field contoured every 100 m. Thin contours show the anomaly field contoured every 10 m, with negative contours dashed and the zero contour omitted.

air from the subtropics have been advected northward, forming a blocking anticyclone (W08).

c. The NAO time scale

To strengthen the case for a regime view of the NAO we now examine the time scales of the positive and negative phases for evidence that these are different, which may indicate the presence of differing dynamics. The decay of positive and negative phases was defined as follows. First, a threshold of plus or minus one standard deviation of the NAO index was used to identify positive and negative NAO days. Then for each of these days the number of subsequent days that were of the same phase was counted. The data were then sorted to give a count of the number of events that last at least n days. To assess significance the decay curves for the NAO events are compared with a simple AR1 model designed to model the NAO index. The two model parameters (variance and lag−1 autocorrelation) were calculated from the daily NAO index for each winter and then averaged over all winters. Following Keeley et al. (2009), the lag−1 autocorrelation was estimated from an exponential fitted to the autocorrelation function at a lag of 5 days.

The resulting decay curves, and the spread of 1000 AR1 model runs, are shown in Fig. 9. The negative NAO curve shows an enhanced occurrence of events lasting around 10 days when compared to both the positive NAO and to the AR1 model. This 10-day time scale is very similar to the average length of Greenland blocking events identified in W08, and the enhanced persistence on this time scale compared to an AR1 model is a fundamental characteristic of blocking (Masato et al. 2009). This suggests that the two phases of the NAO do have intrinsically different decay characteristics. While the decay of positive NAO events is only slightly outside the range of a red noise process, negative NAO events show enhanced persistence on the time scale associated with blocking. Similar results were obtained by Blessing et al. (2005), Jia et al. (2007), and also recently by Barnes and Hartmann (2010), who attribute the enhanced persistence of the negative NAO phase to enhanced positive eddy feedback when compared to the positive phase.

The significance of the difference in decay time scales between the two phases has been tested by comparison with the AR1 runs. Dots in Fig. 9 mark points where the absolute difference between the two phases is larger than that seen in 95% of the AR1 runs. In addition, in only 2% of the AR1 runs is this pointwise significance test passed at as many points as in the observed data. (In this field significance test, only lags of up to 20 days are used to avoid contamination by the poorly sampled extremely long events.) In this test the difference between the two phases of the NAO is highly significant, though this does of course depend on the definition of the AR1 model. For example, if the observed lag−1 autocorrelation is used directly, the AR1 model shows higher persistence and the difference between the two phases is not significant. However, the procedure of Keeley et al. (2009) is preferred because the autocorrelation function of the NAO does not decay exponentially at lags of 1–2 days, so basing the AR1 model on this autocorrelation is not realistic.

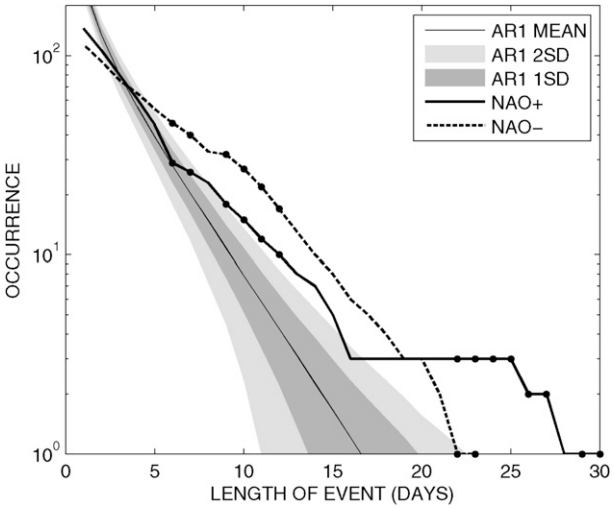

FIG. 9. Semilogarithmic decay curves of the positive and negative NAO phases, expressed as the number of occurrences of events lasting at least n days. Shading marks the one and two standard deviation spread of a first-order Markov (AR1) model. Dots mark points where the difference between positive and negative phases is greater than that in 95% of the AR1 runs.

d. The NAO trend

Long-term changes in the NAO index could arise due to changes in the relative occurrence, or loading, of the two regimes or because of changes in the regimes themselves. Figure 10 shows the full NAO distribution and those of the two regimes from the wave-breaking index, split into the first and second halves of the ERA-40 period.[3] Between the two periods there is a clear shift of the total distribution and also shifts of and changes in population of the two regimes. This suggests that the regimes themselves have changed, with the center of each regime shifted to higher values of the NAO index in the later period. (However, note that, if the mean

[3] This is done using two 22-winter periods, neglecting the final winter of ERA-40. Note that we are not suggesting that this break point has any particular significance; it is simply a convenient way to study the trend over the ERA-40 period.

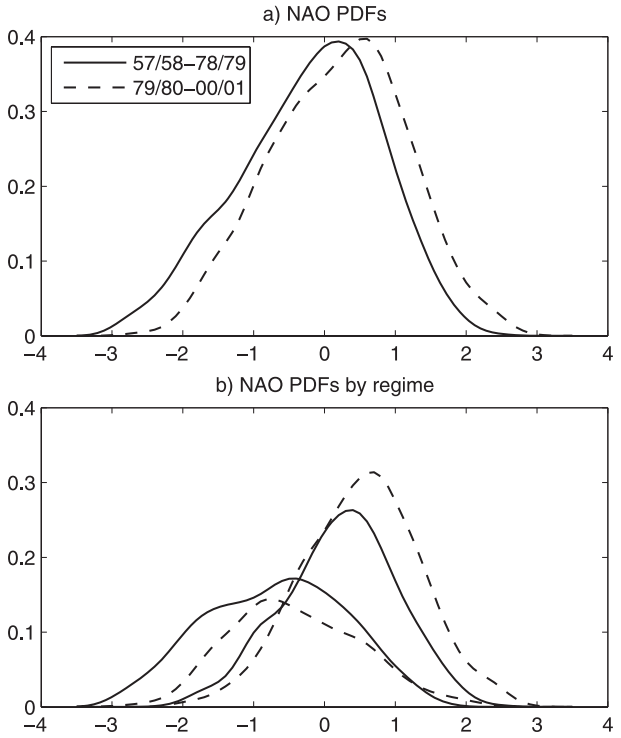

FIG. 10. (a) PDFs of the daily NAO index as in Fig. 3 for the first and second halves of ERA-40. (b) The same PDFs split into Greenland blocking and subpolar jet days for the first (solid lines) and second (dashed lines) periods.

change between the two periods is removed, the change in the total distribution is no longer significant, using a two-sided Kolmogorov–Smirnov test with data sampled every 7 days.)

If composites of Greenland blocking days in the two halves of ERA-40 are made separately, the difference between the two composites projects strongly onto the NAO pattern and is, in fact, very similar to the pattern of total change in winter-mean Z_{500} between the two periods (not shown). This means that the change in the regimes can be summarized by their locations on the linear NAO axis in phase space, reflecting a change in amplitude of the NAO pattern. Figure 11 shows the results of the mixture model applied separately to the first and second halves of ERA-40. Between the two periods the regimes change location and loading in a similar way to that seen in Fig. 10, showing that the two methods of analysis agree well on the changes.

We now try to quantify the contributions of the changes in regime location and loading to the total NAO change between the two periods, using both the wave-breaking index and the mixture model. Here we examine the change over the entire ERA-40 period (cf. Cohen and Barlow 2005) as a basis for comparison with the NAO response to anthropogenic forcing in the GCM analyzed

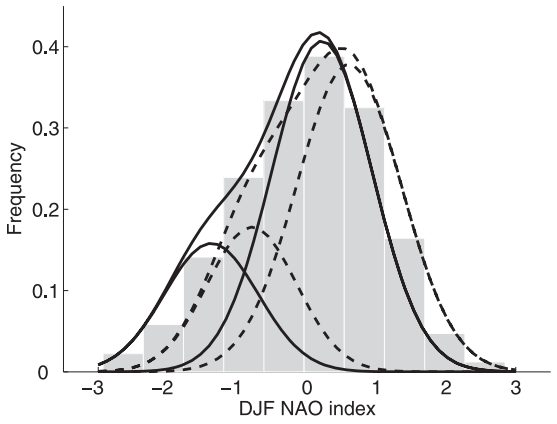

FIG. 11. As in Fig. 7, but showing the results of the two-component mixture model applied separately to the winters 1957/58–1978/79 (solid lines) and 1979/80–2000/01 (dashed lines).

in section 4. Estimated changes in the NAO index I can be calculated from

$$\Delta I_{\text{loading}} = \frac{[(n_{\text{GB2}}\overline{I_{\text{GB}}} + n_{\text{SPJ2}}\overline{I_{\text{SPJ}}}) - (n_{\text{GB1}}\overline{I_{\text{GB}}} + n_{\text{SPJ1}}\overline{I_{\text{SPJ}}})]}{N} \tag{5}$$

and

$$\Delta I_{\text{location}} = \frac{[(\overline{n_{\text{GB}}}I_{\text{GB2}} + \overline{n_{\text{SPJ}}}I_{\text{SPJ2}}) - (\overline{n_{\text{GB}}}I_{\text{GB1}} + \overline{n_{\text{SPJ}}}I_{\text{SPJ1}})]}{N}, \tag{6}$$

where n is the number of days in a given subset, identified by the subscripts GB and SPJ for the Greenland blocking and subpolar jet regimes, and the subscripts 1 and 2 for the first and second periods. An overbar signifies a mean value over the complete time series; N is the total number of days in each period, in this case 1980. Here $\Delta I_{\text{loading}}$ estimates the change in the NAO if the regimes themselves were unchanged but the loading changed as observed, and similarly $\Delta I_{\text{location}}$ estimates the effect of changing only the location of the regimes in NAO space.

In contrast to the wave-breaking index, the mixture model does not assign a classification to each day, but for every value of the NAO index, and hence for every day, it does give the probability of being in each regime. This probability is calculated from the equations for each Gaussian component [Eq. (3)]. These probabilities can be integrated to give the expected number of days of each regime in each period. The parameter values needed to evaluate the two formulas, and the resulting values of $\Delta I_{\text{loading}}$ and $\Delta I_{\text{location}}$, are given in Table 2. The

TABLE 2. Values of parameters defined in section 3d for both methods.

	Wave-breaking index	Mixture model
$\overline{n_{GB}}$	804	611.5
$\overline{n_{SPJ}}$	1176	1368.5
$\overline{I_{GB}}$	−0.56	−0.98
$\overline{I_{SPJ}}$	0.38	0.43
n_{GB1}	934	720
n_{SPJ1}	1046	1260
I_{GB1}	−0.68	−1.10
I_{SPJ1}	0.21	0.30
n_{GB2}	674	503
n_{SPJ2}	1306	1477
I_{GB2}	−0.40	−0.81
I_{SPJ2}	0.51	0.55
$\Delta I_{loading}$	0.12 (29%)	0.15 (37%)
$\Delta I_{location}$	0.29 (71%)	0.26 (63%)

total change in the NAO index is 0.41, so the two estimated contributions appear to combine linearly. Both regime partitions give similar results for the partition: 60%–70% of the change in the NAO index is due to changes in the location of the regimes and the remaining 30%–40% is due to changes in the loading of the regimes.

However, while this method gives similar results for both partitions, other approaches yield different answers. For example, if we use the partition given by the wave-breaking index, we can estimate the contribution of the change in regime loading to the linear NAO trend over the period. To do this we define an index of the occurrence of the Greenland blocking regime to be a simple count of the number of Greenland blocking days in any 90-day winter (as in W08; see Fig. 11). A winter NAO index is derived by simply averaging the daily index over each winter; then linear regression is used to remove the part of NAO variability that is linearly associated with the Greenland blocking time series. By this procedure for removing the changes in loading of the regimes, the linear trend of the NAO is reduced by 65% and the difference of the mean NAO index between the first and second halves is reduced by 53%. Given that different methods give differing results, all that can be concluded about the NAO trend is that contributions from changes in both regime loading and location were important and of the same order.

The contribution of changes in both regime loading and location to the trend in the NAO index contrasts with the hypothesis that the atmospheric response to forcing would be felt as a change in regime loading only, with the regime structures remaining relatively stable (Palmer 1999). To test this we performed a best fit test, fitting the regime PDFs derived from the full ERA-40 period to the NAO distribution in the two subperiods. The question is whether the best fit to the NAO distributions of the two periods is obtained by (A) varying the loading of the two regimes while keeping their location fixed or (B) varying their location while keeping the loading fixed. This test was performed both for the regimes defined by the wave-breaking index and those defined by the mixture model. In both cases and for both periods the rms difference between the PDFs was over twice as large in test A as in test B. Changing the location of the regimes therefore gives a better fit to the distribution in the two subperiods than changing the loading of the regimes. This suggests that the change in the NAO over the ERA-40 period is not consistent with the hypothesis of a change in regime loading only.

4. The NAO in HadCM3

In this section we analyze the NAO distribution in the two model runs. Figure 12 shows the first Atlantic EOFs of monthly mean Z_{500} anomalies from the control run, the doubled CO_2 run, and also the combined set of anomalies from both runs. There are only small differences between the three patterns. For example, the southern NAO center is located slightly farther to the east and extends deeper into Europe in the doubled CO_2 run than in the control run. The change in pattern is smaller than that seen by Ulbrich and Christoph (1999) in a different coupled model. Here we use the pattern from the combined data (the right-hand plot in Fig. 12) as the model's Z_{500} NAO pattern. The same process was used to derive an MSLP pattern for the NAO (not shown). As for Z_{500}, the difference in pattern between the two runs is very small, so the pattern obtained from the combined dataset was used.

To derive daily NAO indices for the runs we projected the daily Z_{500} anomalies onto the NAO pattern, as for ERA-40. The resulting (centered) NAO distributions are shown in Fig. 13, along with the PDFs estimated by the mixture model. The NAO index in the control run does not have the clear skewness seen in the observed NAO (the skewness is only −0.07). The mixture model does identify two significant regimes in both simulations, as shown by the mixing proportions in Fig. 14, although the second regime is only just significant. For the control run the two regimes have very similar populations, consistent with the small level of skewness but the presence of negative kurtosis.[4] The kurtosis is clear from the QQ plot in Fig. 15, as both tails of the NAO distribution are

[4] The kurtosis is −0.49, defined here as the difference from the value 3, which corresponds to a normal distribution.

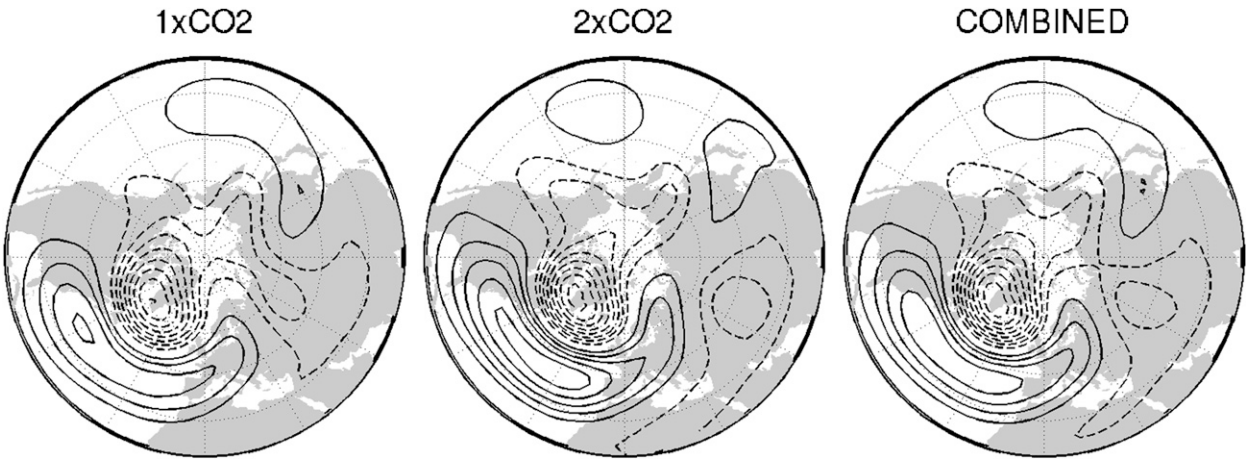

FIG. 12. Maps of the Z_{500} NAO pattern from HadCM3, defined as the first EOF of Z_{500} over 20°–90°N, 90°W–90°E. Patterns are derived from (left) the control run, (middle) the doubled CO_2 run, and (right) the combined dataset of both runs.

underpopulated compared to a standard normal. Interestingly, the NAO in the doubled CO_2 run does have similar skewness to that in observations, with two quite significant regimes of unequal loading.

The model runs are each 100 years long, but we only have 45 years of observations, so before concluding that the model and observed NAO distributions are different we first test whether the skewness of the observed distribution is outside the range seen in 45-yr periods of the model run. In 45-yr periods the model's NAO skewness is never stronger than -0.16, so it appears that the model is significantly different from ERA-40 in this regard, though it would be desirable to have a longer control run to increase confidence in this result.

Two conclusions can be drawn from these results. First, the symmetric nature of the NAO distribution in the control run shows that the model does have deficiencies in its simulation of the NAO when compared to observations. Second, there does appear to be a change in the loading of the two regimes under anthropogenic forcing, with the positive NAO regime becoming more dominant.

To directly compare the NAO in the two runs, the PDFs of the NAO (as estimated by the kernel method) are shown together in Fig. 16. Atlantic EOF-based indices of the NAO using both Z_{500} and MSLP are shown here. The change in the shape of the distribution under anthropogenic forcing is the same at both levels. However, there is a clear difference between the two levels in that under greenhouse forcing the mean NAO index increases at the surface but decreases at upper levels. (as described in section 2b, all daily anomalies are calculated as anomalies from the seasonal cycle of the control run in order to preserve this mean change).

The mean changes in Z_{500} and MSLP are shown in Fig. 17. This also shows opposite behavior at the two levels. At the surface the response to forcing is a decrease in pressure at high latitudes and a (small) increase in the subtropics, while at 500 hPa the height increases everywhere, but with a larger increase at high latitudes. Neither of the responses project particularly strongly onto the NAO pattern. In short, the response is very baroclinic, having opposite signs at the two levels. This response is seen in many other climate models, as described by Woollings (2008), and simply reflects the uneven distribution of warming that affects the height field through hydrostatic balance.

The baroclinic nature of the circulation change means that, in general, it is not easy to diagnose changes in the NAO. However, the similarity of the change in shape of the NAO distribution at the two levels (Fig. 16) does reveal that the NAO itself is changed, with the positive phase becoming more dominant. At both levels the change in the shape of the distribution is highly significant. This was assessed using a Monte Carlo method, pooling 10-day means of the data from both simulations and splitting them randomly into two samples. At 500 hPa the change in skewness between the two simulations was only reproduced in 1% of 5000 trials, while at the surface it was reproduced in only 0.5% of the trials.

It is instructive to directly compare the control simulation with the ERA-40 analysis data, even though this simulation uses preindustrial, rather than present day, emissions. Figure 18 shows the difference in winter-mean Z_{500}. The model has lower heights on average due to the different radiative forcing but also a strong pattern of bias that projects very strongly onto the negative NAO and is also similar in the western Pacific. Comparison to Fig. 17 shows that with respect to geopotential height gradients the model bias is much larger than the difference between the two model runs. If the pattern in

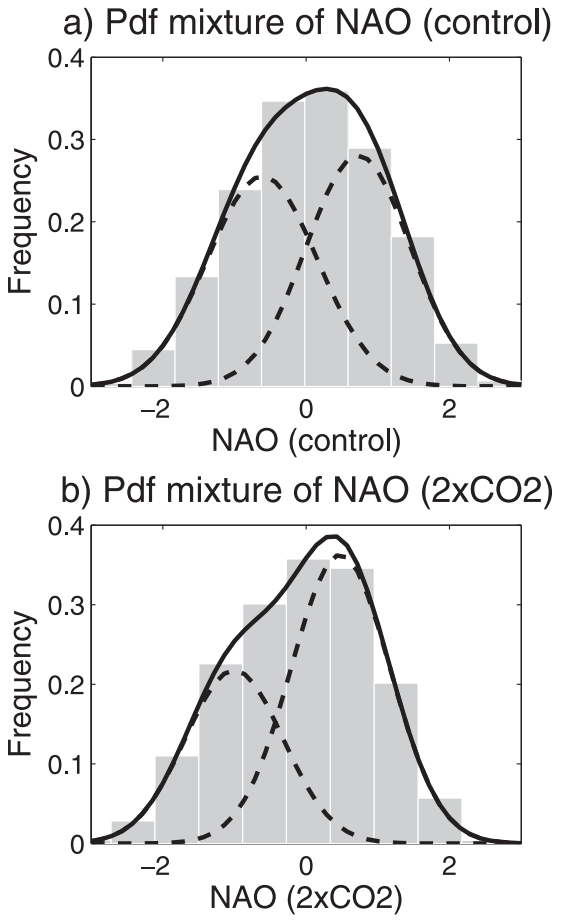

FIG. 13. Distributions of the daily Z_{500} NAO index from the HadCM3 runs, along with the PDFs of the mixture model fit (solid line) and the individual mixture model components (dashed lines) for (a) the control run and (b) the doubled CO_2 run.

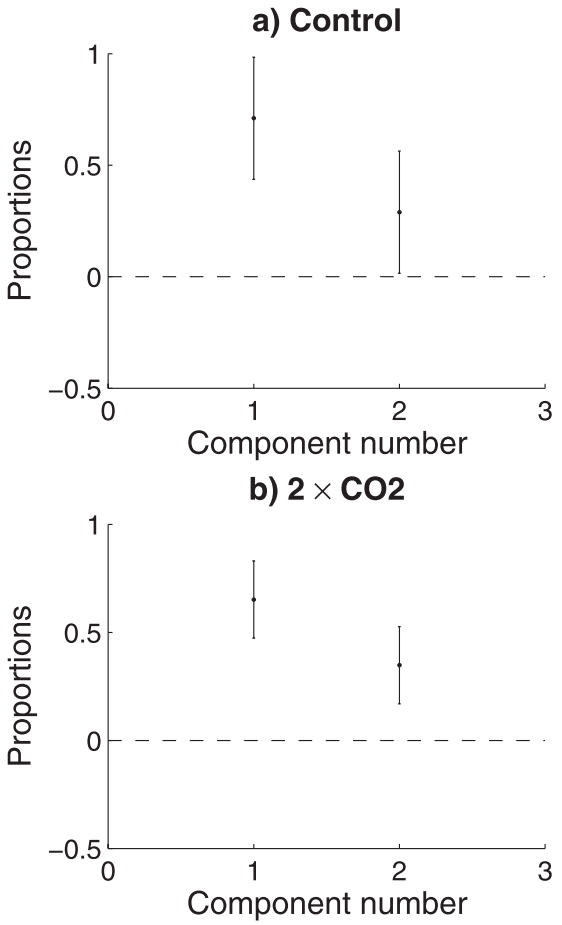

FIG. 14. The component proportions α for the two-regime fits in Fig. 13 along with their 95% confidence intervals.

Fig. 18 is projected onto the observed NAO pattern, the bias equates to a NAO index of −0.69. This strong bias is not just a feature of the 30-level version of the model. We have verified it using a control run of the conventional 19-level model and it also emerges in the analysis of Stephenson et al. (2006; see Fig. 4). This strong bias is consistent with the results of the mixture model that the loading of the Greenland blocking regime is too high, and also casts doubt over the model's ability to predict future changes.

Finally, we examine the decay rate of NAO events in the model, for comparison with the observed behavior described in section 3c. Figure 19 shows the decay curves for both model runs.[5] In contrast to ERA-40, the positive and negative NAO decay curves for the control run are similar to each other, so in this regard the asymmetry of the NAO is not well simulated in the model. In the doubled CO_2 run the decay curves appear slightly different, although this difference is mostly in the relatively small number of very long lasting events. This contrasts with the difference between positive and negative NAO events in ERA-40, which differ for events lasting around 10 days, that is, the typical lifetime of blocking.

5. Conclusions

We have shown that the distribution of the winter NAO index is negatively skewed and that this is robust to choices of the flow variable, vertical level, time scale, and method used to derive the index. In fact, Christiansen (2009) identified the NAO as the atmospheric circulation pattern most associated with skewness. We have also shown that there is a clear physical difference between positive and negative NAO phases, in that they have different decay time scales. Negative NAO events

[5] Note that the NAO distribution for the doubled CO_2 run was normalized before performing this analysis, as otherwise it would not be centered.

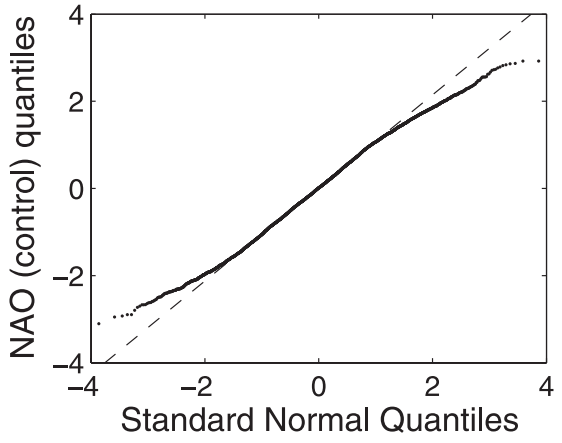

FIG. 15. QQ plot of the NAO in the HadCM3 control run against a standard normal.

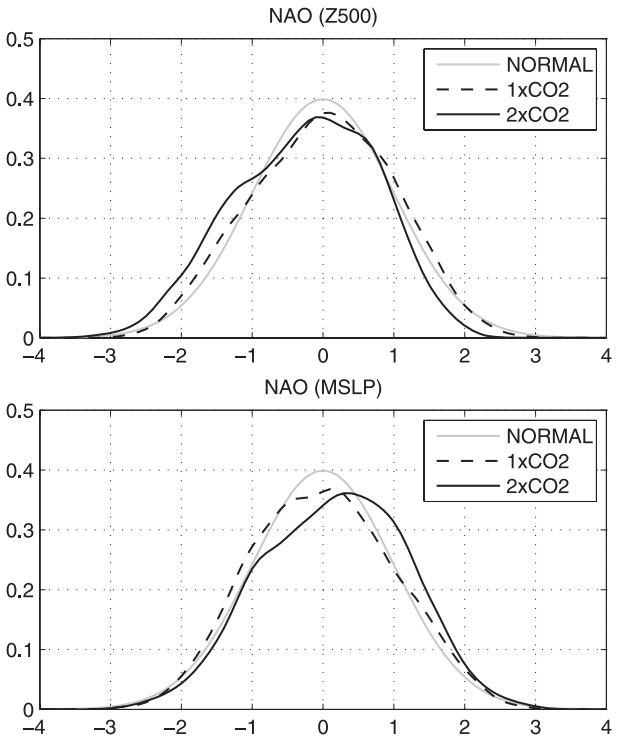

FIG. 16. Distributions of the NAO in the control and doubled CO_2 runs of HadCM3 in (top) Z_{500} and (bottom) MSLP, as estimated by the kernel method. A normal distribution is shown for comparison.

exhibit enhanced persistence on the 10-day time scale when compared to positive NAO events, consistent with the concept of blocking.

We suggest that the skewness of the NAO reflects the existence of a distinct flow regime in the negative NAO phase that, following W08, we call Greenland blocking. This interpretation is supported by both a dynamical wave-breaking index and a statistical mixture model. Roughly 30%–40% of all days are classed as Greenland blocking, depending on the method of classification, and the rest are simply referred to as subpolar jet days, although presumably there may be several intrinsically different situations within this set. We have only searched for regimes along the NAO axis in phase space, so regimes resembling the eastern Atlantic pattern, for example, will not be seen here. However, they may influence our results to the extent that they project onto the NAO pattern. For example, it seems likely that some Scandinavian blocking days (Pelly and Hoskins 2003; Tyrlis and Hoskins 2008) contribute to the extension of the Greenland blocking PDF into positive NAO values, as discussed in section 3b.

There is, of course, much debate over whether the existence of preferred flow regimes is significant. We have assessed this via the significance criteria embedded within the mixture model, and by this measure the regimes are, indeed, significant. The combination of physical and statistical analyses also adds confidence when compared to results obtained using purely statistical methods. One clear difference between this study and many others is that by focusing on the NAO we have ensured that we search for local regimes so that there is no danger of mixing independent variations from different sectors, as may happen in studies using hemispheric EOFs for example (Stephenson et al. 2004).

There are other possible interpretations of the skewness of the NAO. For example, the work of Rennert and Wallace (2009) suggests that flow skewness in this region arises from the cross-frequency coupling of disturbances with long and intermediate time scales. In this interpretation the NAO is seen as an intrinsically long time scale phenomenon, and the intermediate time scales are dominated by retrograding long Rossby waves. These two phenomena are shown to interfere constructively during the negative NAO phase and destructively during the positive phase, which leads to the skewness.

The NAO exhibited a well-known positive trend over the ERA-40 period, and it is of interest to determine the extent to which such long time scale changes are due to changes in the loading of the two regimes. From the analysis presented in section 3d, it seems that such a change in regime loading did indeed contribute to the observed trend. However, the locations of both regimes in NAO space also changed, moving toward more positive values of the NAO index, and this background change also contributed to the trend.

It is clear that the climate model investigated here, HadCM3, has significant deficiencies in its representation of the NAO, in that the model has a strong mean bias, the NAO distribution in the control run is too

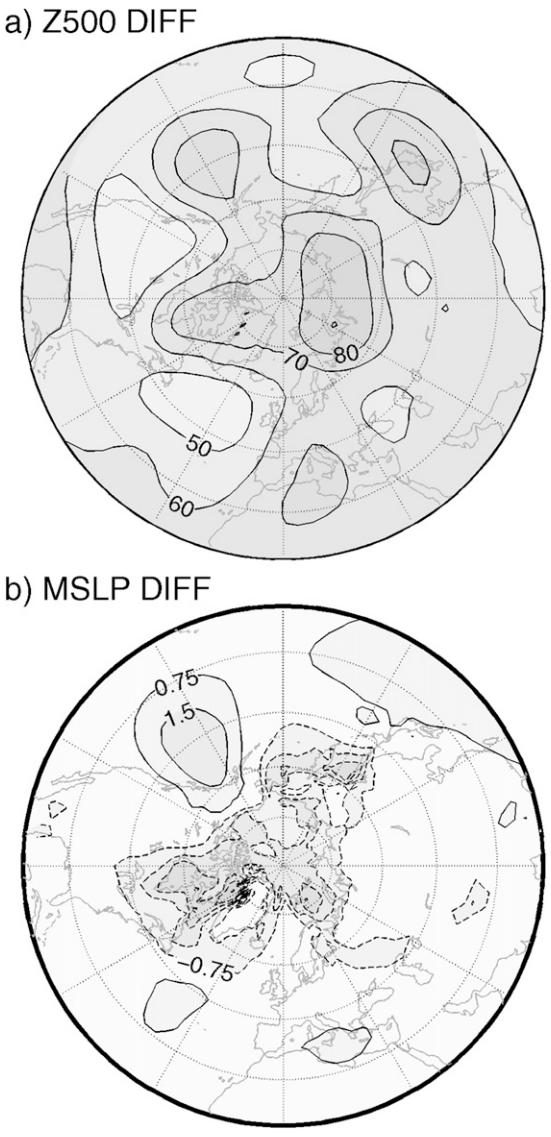

FIG. 17. The difference between winter means of the two HadCM3 runs (doubled CO_2 minus control) for (a) Z_{500}, with contours every 10 m, and (b) MSLP, with contours every 0.75 hPa. In both panels negative contours are dashed and the zero contour omitted.

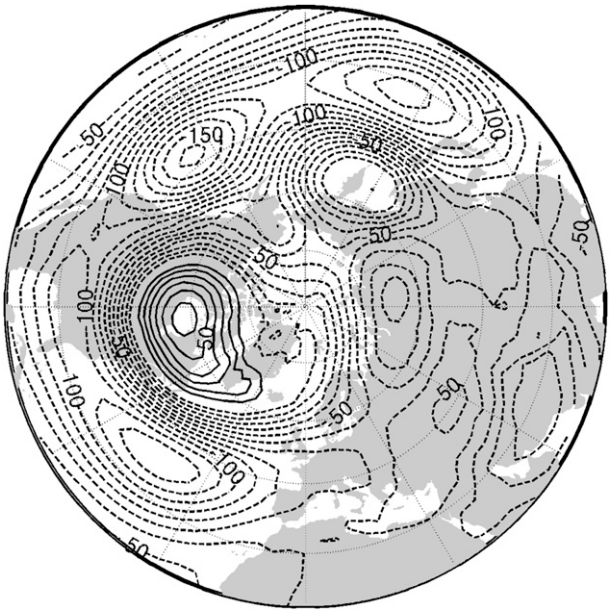

FIG. 18. Difference field of the DJF-mean Z_{500} 1 × CO_2 run ERA-40. Contours are drawn every 10 m, with negative contours dashed and the zero contour omitted.

symmetric, and the decay time scales are unrealistic. Therefore, confidence in the model's projections of NAO change must be considered to be low. However, we have shown that the nature of NAO variability does change under anthropogenic forcing with an increase in the dominance of the subpolar jet regime, and this change is highly significant. The change in shape of the NAO distribution is the same at both the surface and the midtroposphere, so the change in the NAO can be distinguished from the hydrostatic signature of uneven warming. That the nature of NAO variability is changed by the forcing is important, even though the actual change is small in this case. Given the model deficiencies, it is quite possible that this change is an underestimate. The NAO has dominated the low-frequency variability of European climate in recent decades, and changes in it have the potential to significantly modify the pattern of climate change over Europe. It is therefore important to investigate in detail how well the NAO is represented in current climate models and how they predict it will change.

Finally, we note that, if this regime view of the NAO is correct, the recent NAO trend does not seem to be consistent with the hypothesis that dynamical climate change will be felt as a change in the regime loading only, with the regime structures remaining relatively stable. Our analysis interprets the NAO change over the ERA-40 period as being associated with a change in both the loading and the structure of the regimes. Ulbrich and Cristoph (1999), Brandefelt (2006), and Branstator and Selten (2009) give similar examples of changes in regime structure in numerical models. Also, the baroclinic circulation change in the climate model simulations provides a simple example of how circulation changes can project onto patterns of variability without there being a change in the nature of the variability itself.

Acknowledgments. We are indebted to ECMWF for providing the ERA-40 reanalysis data and to Mike Wallace and the anonymous reviewers for their most useful suggestions.

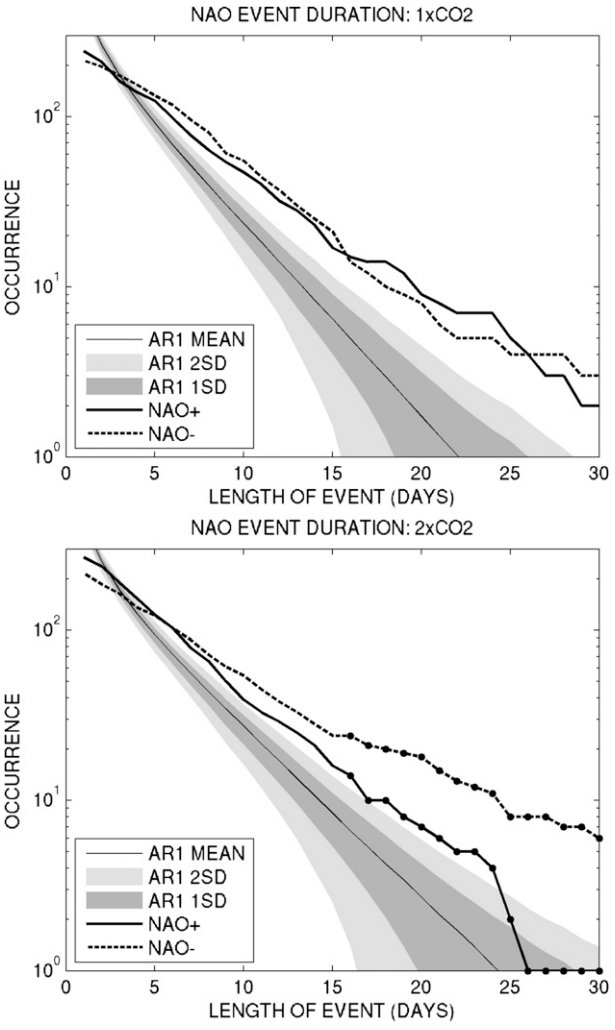

FIG. 19. Decay curves, as in Fig. 9, for the positive and negative NAO events in the two model runs: (top) control run and (bottom) doubled CO_2 run.

REFERENCES

Ambaum, M. H. P., B. J. Hoskins, and D. B. Stephenson, 2001: Arctic Oscillation or North Atlantic Oscillation? *J. Climate,* **14,** 3495–3507.

Anderson, B. D., and J. B. Moore, 1979: *Optimal Filtering*. Prentice-Hall, 357 pp.

Barnes, E. A., and D. L. Hartmann, 2010: Dynamical feedbacks and the persistence of the NAO. *J. Atmos. Sci.,* in press.

Benedict, J. J., S. Lee, and S. B. Feldstein, 2004: Synoptic view of the North Atlantic Oscillation. *J. Atmos. Sci.,* **61,** 121–144.

Berner, J., 2005: Linking nonlinearity and non-Gaussianity of planetary wave behavior by the Fokker–Planck equation. *J. Atmos. Sci.,* **62,** 2098–2117.

Berrisford, P., B. J. Hoskins, and E. Tyrlis, 2007: Blocking and Rossby wave breaking on the dynamical tropopause in the Southern Hemisphere. *J. Atmos. Sci.,* **64,** 2881–2898.

Bjerknes, J., 1964: Atlantic air–sea interaction. *Advances in Geophysics,* Vol. 10, Academic Press, 1–82.

Blessing, S., K. Fraedrich, M. Junge, T. Kunz, and F. Lunkeit, 2005: Daily North Atlantic Oscillation (NAO) index: Statistics and its stratospheric polar vortex dependence. *Meteor. Z.,* **14,** 763–769.

Brandefelt, J., 2006: Atmospheric modes of variability in a changing climate. *J. Climate,* **19,** 5934–5943.

Branstator, G., and J. Berner, 2005: Linear and nonlinear signatures in the planetary wave dynamics of an AGCM: Phase space tendencies. *J. Atmos. Sci.,* **62,** 1792–1811.

——, and F. Selten, 2009: ''Modes of variability'' and climate change. *J. Climate,* **22,** 2639–2658.

Cassou, C., L. Terray, J. W. Hurrell, and C. Deser, 2004: North Atlantic winter climate regimes: Spatial asymmetry, stationarity with time, and oceanic forcing. *J. Climate,* **17,** 1055–1068.

Charney, J. G., and J. G. Devore, 1979: Multiple flow equilibria in the atmosphere and blocking. *J. Atmos. Sci.,* **36,** 1205–1216.

Cheng, X., and J. M. Wallace, 1993: Cluster analysis of the Northern Hemisphere wintertime 500-hPa height field: Spatial patterns. *J. Atmos. Sci.,* **50,** 2674–2696.

Christiansen, B., 2005a: Bimodality of the planetary-scale atmospheric wave amplitude index. *J. Atmos. Sci.,* **62,** 2528–2541.

——, 2005b: The shortcomings of nonlinear principal component analysis in identifying circulation regimes. *J. Climate,* **18,** 4814–4823.

——, 2007: Atmospheric circulation regimes: Can cluster analysis provide the number? *J. Climate,* **20,** 2229–2250.

——, 2009: Is the atmosphere interesting? A projection pursuit study of the circulation in the Northern Hemisphere winter. *J. Climate,* **22,** 1239–1254.

Cohen, J., and M. Barlow, 2005: The NAO, the AO, and global warming: How closely related? *J. Climate,* **18,** 4498–4513.

Coppola, E., F. Kucharski, F. Giorgi, and F. Molteni, 2005: Bimodality of the North Atlantic Oscillation in simulations with greenhouse gas forcing. *Geophys. Res. Lett.,* **32,** L23709, doi:10.1029/2005GL024080.

Corti, S., F. Molteni, and T. N. Palmer, 1999: Signature of recent climate change in frequencies of natural atmospheric circulation regimes. *Nature,* **398,** 799–802.

Croci-Maspoli, M., C. Schwierz, and H. C. Davies, 2007: Atmospheric blocking: Space–time links to the NAO and PNA. *Climate Dyn.,* **29,** 713–725.

Crommelin, D. T., 2004: Observed nondiffusive dynamics in large-scale atmospheric flow. *J. Atmos. Sci.,* **61,** 2384–2396.

Egger, J., 1981: Stochastically driven large-scale circulation with multiple equilibria. *J. Atmos. Sci.,* **38,** 2606–2618.

Everitt, B., and D. J. Hand, 1981: *Finite Mixture Distributions*. Chapman and Hall, 143 pp.

Feldstein, S. B., 2000: The timescale, power spectra, and climate noise properties of teleconnection patterns. *J. Climate,* **13,** 4430–4440.

Franzke, C., S. Lee, and S. B. Feldstein, 2004: Is the North Atlantic Oscillation a breaking wave? *J. Atmos. Sci.,* **61,** 145–160.

——, D. T. Crommelin, A. Fischer, and A. J. Majda, 2008: A hidden Markov model perspective on regimes and metastability in atmospheric flows. *J. Climate,* **21,** 1740–1757.

Haines, K., and A. Hannachi, 1995: Weather regimes in the Pacific from a GCM. *J. Atmos. Sci.,* **52,** 2444–2462.

Handorf, D., K. Dethloff, A. G. Marshall, and A. Lynch, 2009: Climate regime variability for past and present time slices simulated by the Fast Ocean Atmosphere Model. *J. Climate,* **22,** 58–70.

Hannachi, A., 1997: Low-frequency variability in a GCM: Three-dimensional flow regimes and their dynamics. *J. Climate,* **10,** 1357–1379.

——, 2007: Tropospheric planetary wave dynamics and mixture modeling: Two preferred regimes and a regime shift. *J. Atmos. Sci.,* **64,** 3521–3541.

——, and A. O'Neill, 2001: Atmospheric multiple equilibria and non-Gaussian behaviour in model simulations. *Quart. J. Roy. Meteor. Soc.,* **127,** 939–958.

——, and A. G. Turner, 2008: Preferred structures in large-scale circulation and the effect of doubling greenhouse gas concentration in HadCM3. *Quart. J. Roy. Meteor. Soc.,* **134,** 469–480.

——, I. T. Jolliffe, and D. B. Stephenson, 2007: Empirical orthogonal functions and related techniques in atmospheric science: A review. *Int. J. Climatol.,* **27,** 1119–1152.

Hsu, C. J., and F. Zwiers, 2001: Climate change in recurrent regimes and modes of atmospheric variability. *J. Geophys. Res.,* **106** (D17), 20 145–20 160.

Hurrell, J. W., 1995: Decadal trends in the North Atlantic Oscillation: Regional temperatures and precipitation. *Science,* **269,** 676–679.

——, Y. Kushnir, G. Ottersen, and M. Visbeck, 2002: An overview of the North Atlantic Oscillation. *The North Atlantic Oscillation—Climate Significance and Environmental Impact, Geophys. Monogr.,* Vol. 134, Amer. Geophys. Union, 1–35.

Jia, X., J. Derome, and H. Lin, 2007: Comparison of the life cycles of the NAO using different definitions. *J. Climate,* **20,** 5992–6011.

Jones, P. D., T. Jónsson, and D. Wheeler, 1997: Extension to the North Atlantic Oscillation using early instrumental pressure observations from Gibraltar and south-west Iceland. *Int. J. Climatol.,* **17,** 1433–1450.

Keeley, S. P. E., M. Collins, and A. J. Thorpe, 2008: Northern Hemisphere winter atmospheric climate: Modes of natural variability and climate change. *Climate Dyn.,* **31,** 195–211, doi:10.1007/s00382-007-0346-6.

——, R. Sutton, and L. Shaffrey, 2009: Does the North Atlantic Oscillation show unusual persistence on intraseasonal timescales? *Geophys. Res. Lett.,* **36,** L22706, doi:10.1029/2009GL040367.

Kimoto, M., and M. Ghil, 1993: Multiple flow regimes in the Northern Hemisphere winter. Part I: Methodology and hemispheric regimes. *J. Atmos. Sci.,* **50,** 2625–2643.

Legras, B., and M. Ghil, 1985: Persistent anomalies, blocking and variations in atmospheric predictability. *J. Atmos. Sci.,* **42,** 433–471.

Luo, D., T. Gong, and Y. Diao, 2007: Dynamics of eddy-driven low-frequency dipole modes. Part III: Meridional displacement of westerly jet anomalies during two phases of NAO. *J. Atmos. Sci.,* **64,** 3232–3248.

Masato, G., B. J. Hoskins, and T. Woollings, 2009: Can the frequency of blocking be described by a red noise process? *J. Atmos. Sci.,* **66,** 2143–2149.

McLachlan, G. I., and K. E. Basford, 1988: *Mixture Models: Inference and Applications to Clustering.* Marcel Dekker, 253 pp.

Mo, K. C., and M. Ghil, 1987: Statistics and dynamics of persistent anomalies. *J. Atmos. Sci.,* **44,** 877–902.

Monahan, A. H., J. C. Fyfe, and G. M. Flato, 2000: A regime view of Northern Hemisphere atmospheric variability and change under global warming. *Geophys. Res. Lett.,* **27,** 1139–1142.

——, L. Pandolfo, and J. Fyfe, 2001: The preferred structure of variability of the Northern Hemisphere atmospheric circulation. *Geophys. Res. Lett.,* **28,** 1019–1022.

Palmer, T. N., 1999: A nonlinear dynamical perspective on climate prediction. *J. Climate,* **12,** 575–591.

Pelly, J. L., and B. J. Hoskins, 2003: A new perspective on blocking. *J. Atmos. Sci.,* **60,** 743–755.

Rennert, K. J., and J. M. Wallace, 2009: Cross-frequency coupling, skewness, and blocking in the Northern Hemisphere winter circulation. *J. Climate,* **22,** 5650–5666.

Silverman, B. W., 1981: Using kernel density estimates to investigate multimodality. *J. Roy. Stat. Soc.,* **43B,** 97–99.

Smyth, P., M. Ghil, and K. Ide, 1999: Multiple regimes in Northern Hemisphere height fields via mixture model clustering. *J. Atmos. Sci.,* **56,** 3704–3723.

Stephenson, D. B., V. Pavan, and R. Bojariu, 2000: Is the North Atlantic Oscillation a random walk? *Int. J. Climatol.,* **20,** 1–18.

——, A. Hannachi, and A. O'Neill, 2004: On the existence of multiple climate regimes. *Quart. J. Roy. Meteor. Soc.,* **130,** 583–605.

——, V. Pavan, M. Collins, M. M. Junge, and R. Quadrelli, 2006: North Atlantic Oscillation response to transient greenhouse gas forcing and the impact on European winter climate: A CMIP2 multi-model assessment. *Climate Dyn.,* **27,** 401–420.

Straus, D. M., S. Corti, and F. Molteni, 2007: Circulation regimes: Chaotic variability versus SST-forced predictability. *J. Climate,* **20,** 2251–2272.

Strong, C., and G. Magnusdottir, 2008: Tropospheric Rossby wave breaking and the NAO/NAM. *J. Atmos. Sci.,* **65,** 2861–2876.

Sura, P., M. Newman, C. Penland, and P. Sardeshmukh, 2005: Multiplicative noise and non-Gaussianity: A paradigm for atmospheric regimes? *J. Atmos. Sci.,* **62,** 1391–1409.

Teng, Q., A. H. Monahan, and J. C. Fyfe, 2004: Effects of time averaging on climate regimes. *Geophys. Res. Lett.,* **31,** L22203, doi:10.1029/2004GL020840.

Terray, L., M. E. Demory, M. Déqué, G. de Coetlogon, and E. Maisonnave, 2004: Simulation of late twenty-first century changes in wintertime atmospheric circulation over Europe due to anthropogenic causes. *J. Climate,* **17,** 4630–4635.

Thompson, D. W. J., S. Lee, and M. P. Baldwin, 2002: Atmospheric processes governing the Northern Hemisphere Annular Mode/North Atlantic Oscillation. *The North Atlantic Oscillation—Climate Significance and Environmental Impact, Geophys. Monogr.,* Vol. 134, Amer. Geophys. Union, 81–112.

Tyrlis, E., and B. J. Hoskins, 2008: The morphology of Northern Hemisphere blocking. *J. Atmos. Sci.,* **65,** 1653–1665.

Ulbrich, U., and M. Christoph, 1999: A shift of the NAO and increasing storm track activity over Europe due to anthropogenic greenhouse gas forcing. *Climate Dyn.,* **15,** 551–559.

Uppala, S. M., and Coauthors, 2005: The ERA-40 Re-Analysis. *Quart. J. Roy. Meteor. Soc.,* **131,** 2961–3012.

Vallis, G. K., and E. P. Gerber, 2008: Local and hemispheric dynamics of the North Atlantic Oscillation, annular patterns and the zonal index. *Dyn. Atmos. Oceans,* **44,** 184–212.

Vautard, R., 1990: Multiple weather regimes over the North Atlantic: Analysis of precursors and successors. *Mon. Wea. Rev.,* **118,** 2056–2081.

Wallace, J. M., X. Cheng, and D. Sun, 1991: Does low-frequency atmospheric variability exhibit regime-like behavior? *Tellus,* **43A,** 16–26.

Wanner, H., S. Brönnimann, C. Casty, D. Gyalistras, J. Luterbacher, C. Schmutz, D. B. Stephenson, and E. Xoplaki, 2001: North Atlantic Oscillation: Concepts and studies. *Survey Geophys.,* **22,** 321–381.

Woollings, T., 2008: The vertical structure of anthropogenic zonal-mean atmospheric circulation change. *Geophys. Res. Lett.,* **35,** L19702, doi:10.1029/2008GL034883.

——, B. Hoskins, M. Blackburn, and P. Berrisford, 2008: A new Rossby wave–breaking interpretation of the North Atlantic Oscillation. *J. Atmos. Sci.,* **65,** 609–626.