

Predictability and Non-Gaussian Characteristics of the North Atlantic Oscillation

THOMAS ÖNSKOG

Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden

CHRISTIAN L. E. FRANZKE

Meteorological Institute, and Center for Earth System Research and Sustainability, University of Hamburg, Hamburg, Germany

ABDEL HANNACHI

Department of Meteorology, Stockholm University, Stockholm, Sweden

(Manuscript received 17 February 2017, in final form 7 September 2017)

ABSTRACT

The North Atlantic Oscillation (NAO) is the dominant mode of climate variability over the North Atlantic basin and has a significant impact on seasonal climate and surface weather conditions. It is the result of complex and nonlinear interactions between many spatiotemporal scales. Here, the authors study the statistical properties of two time series of the daily NAO index. Previous NAO modeling attempts only considered Gaussian noise, which can be inconsistent with the system complexity. Here, it is found that an autoregressive model with non-Gaussian noise provides a better fit to the time series. This result holds also when considering time series for the four seasons separately. The usefulness of the proposed model is evaluated by means of an investigation of its forecast skill.

1. Introduction


a. Background

The atmospheric circulation is highly complex and involves a large number of variables and varies on a wide range of spatiotemporal scales. The large-scale flow steers to a large extent the synoptic-scale and daily weather systems. This large-scale flow can be described by a relatively small number of persistent flow patterns referred to as teleconnections (Wallace and Gutzler 1981; Barnston and Livezey 1987; Feldstein and Franzke 2017; Hannachi et al. 2017), which have a significant influence on medium- and longer-term weather and climate prediction (e.g., Colucci and Baumhefner 1992). Examples of such teleconnection patterns include the North Atlantic Oscillation (NAO), the Pacific–North American (PNA) pattern, and the Scandinavian (SCA) pattern, which control substantial parts of the atmospheric circulation

over the northern Atlantic, North America, and Scandinavia, respectively (Feldstein and Franzke 2017).

A characteristic feature of the large-scale atmospheric flow structure and the teleconnections is their non-Gaussianity (e.g., Proistosescu et al. 2016; Sura and Hannachi 2015; Hannachi 2010), and various mechanisms that could explain the non-Gaussian statistics of atmospheric synoptic and low-frequency variability have been proposed in the literature (e.g., Sura and Hannachi 2015). A proper understanding and modeling of this characteristic feature is invaluable and has important consequences, particularly in climate research, as weather and climate risk assessment depends on extreme events, which are reflected in the non-Gaussian behavior of the large-scale atmospheric flow structure (Sura and Hannachi 2015; Proistosescu et al. 2016; Franzke 2017).

As mentioned above, the NAO (Hurrell et al. 2003) is one of the main Northern Hemispheric teleconnection patterns and controls much of the atmospheric variability, including westerly winds, surface temperature, precipitation, and surface pressure, over the North Atlantic basin, the eastern United States, and Europe. The NAO represents the dominant mode of sea level pressure variability over the North Atlantic Ocean and Europe and

 Denotes content that is immediately available upon publication as open access.

Corresponding author: Thomas Önskog, onskog@kth.se

DOI: 10.1175/JCLI-D-17-0101.1

© 2018 American Meteorological Society. For information regarding reuse of this content and general copyright information, consult the [AMS Copyright Policy](https://www.ametsoc.org/PUBSReuseLicenses) (www.ametsoc.org/PUBSReuseLicenses).

has a seesaw structure in atmospheric mass between the two centers of action, namely, the Icelandic low and the Azores high. The NAO has irregular behavior but is known to have extended periods of positive and negative phases (e.g., Jung et al. 2011; Hannachi and Stendel 2016) associated with a wide range of time scales from days to decades and longer (Woollings et al. 2015) manifesting interaction with surface conditions including sea surface temperature and sea ice (Stendel et al. 2016). The NAO seems to have been known since the days of the Scandinavian sailors, and from the mid-eighteenth century Egede (1745) and Cranz (1765) noted that surface air temperature at Greenland and Scandinavia vary in opposite directions (Stephenson et al. 2003; Pinto and Raible 2012; Feldstein and Franzke 2017).

The positive phase of the NAO yields stronger westerlies leading to wetter and milder conditions over Europe and the Mediterranean region, whereas the opposite phase is associated with dryer Mediterranean conditions and harsh winters over Europe (Stendel et al. 2016; Hannachi and Stendel 2016). The maps in Fig. 1 (top) illustrate the spatial patterns of the positive and negative NAO phases, respectively. The positive phase shows strong low and high pressure systems leading in particular to strong westerlies oriented more to the northeast toward the British Isles and Scandinavia. This, in turn, induces warm and wet weather conditions over those regions but dry weather conditions over the Mediterranean region and a reduced high pressure over Siberia. There is also an impact on the ocean surface, where a correlation is observed with a northward surface warm current over the North Atlantic. The negative phase is characterized by weaker westerlies, associated with weaker low and high pressure systems, zonally oriented, which lead to wet weather conditions over the Mediterranean region and cold and dry weather over the northern latitudes (e.g., Scandinavia). The high pressure system over Iberia is strengthened and the surface currents are weaker and cooler than during the positive phase (Wanner et al. 2001).

Various interpretations of the NAO exist in the literature. Woollings et al. (2010b), for example, suggest that the NAO can be interpreted in terms of two regimes, namely, a high-latitude (or Greenland) blocking regime and a zonal regime. They suggest, in particular, that changes in the relative occurrence frequency and the structure of the regimes contribute to the long-term (e.g., decadal) NAO trend over the European reanalysis period from ERA-40 (Uppala et al. 2005). Those regime states are strongly linked to Rossby wave breaking (Benedict et al. 2004; Franzke et al. 2004; Woollings et al. 2008; Franzke et al. 2011). Woollings et al. (2010a) and Hannachi et al. (2012) linked the NAO to the (eddy-driven) jet stream variability. The NAO has gone through

phase changes over the instrumental period in the last 200 years; see the discussion in Stendel et al. (2016).

b. Rationale

The aggregation of independent error increments, properly scaled, is known to yield Gaussian statistics. The climate system, however, is characterized by complex nonlinear interactions across spatiotemporal scales, marked by extreme events, and where the different types of forcing, ranging from very small scale (e.g., radiation) to large scale (e.g., orography), play an important role. The interaction with these scales affects the structure of the aggregated weather and climate disturbances. These processes altogether contribute to the departure of the statistics from Gaussianity and can generate heavy-tailed distributions characterized by nonnegligible probability of extremes (Majda et al. 2009; Sardeshmukh and Sura 2009; Sura and Hannachi 2015) and long-range dependence (Franzke et al. 2015). It can be shown that the nonlinear interaction between the slow, here the NAO-associated, circulation with the fast synoptic-scale eddies lead to the emergence of deviations of the circulation statistics from Gaussianity (Majda et al. 2009; Franzke et al. 2005; Sardeshmukh and Sura 2009; Sura and Hannachi 2015). For instance, Benedict et al. (2004), Franzke et al. (2004), and Feldstein and Franzke (2006) have highlighted the importance of the nonlinear interactions in the life cycle of the NAO. It is these nonlinear processes, integral to the NAO evolution, which also cause its extreme states, which we will model using non-Gaussian statistics.

A large class of heavy-tailed distributions exist in the literature, including the t -distribution and α -stable distributions. The former, in particular, is akin to the familiar normal distribution. The use of these distributions can benefit our understanding in modeling weather and climate variables and can improve their predictability and forecast skill. Here we propose a new statistical model for the NAO that accounts also for its observed non-Gaussian characteristics. Non-Gaussian models are widely used in other areas of science (e.g., Franzke 2017; Graves et al. 2017). So far the non-Gaussian features of the NAO have not received much attention, and here we will examine their importance for the predictive skill of the NAO.

The predictability of the NAO is an important research topic. For instance, the National Oceanic and Atmospheric Administration (NOAA) Climate Prediction Center now routinely forecasts the daily state of the most important teleconnection patterns, including the NAO.¹

¹ See http://www.cpc.ncep.noaa.gov/products/precip/CWlink/pna/nao_index_mrf.shtml.

NOAA also performs probabilistic analog 8–14-day forecasts based on the persistence of teleconnection patterns (e.g., NAO).² The predictability of the NAO has been studied by Franzke and Woollings (2011). They find that with respect to seasonal and longer time-scale predictability, the potential predictability is about 52%–57% of the total interannual variance in winter and 29%–30% in summer. The remainder is likely due to climate noise (Franzke 2009) and, thus, is not predictable.

Furthermore, extreme states of the NAO have also severe impacts. For instance, in 2010 the NAO exhibited one of its most negative index states. This was accompanied by one of the coldest winters in the United Kingdom for over 100 years (Moore and Renfrew 2012). It is also well known that the NAO affects the propensity of extreme events to occur (e.g., Hannachi et al. 2017). Thus, understanding and the ability to model extreme states of the NAO index are important. Here we will show that the NAO index has heavy tails, by which we mean that the tails of the NAO decay more slowly than the tails of a corresponding Gaussian distribution. Consequently, extreme events are more likely than for a corresponding Gaussian distribution. Hence, we will also examine which non-Gaussian heavy-tailed distribution provides the best fit to the data and the best predictive skill.

The objective of this article is to study the properties of time series of the daily NAO index and to set up mathematical models for the NAO index based on these observed properties. The main focus will lie on a station-based time series derived from observations of sea level pressure (Cropper et al. 2015), but a time series based on empirical orthogonal functions (EOF) will be considered as well. Because of the observed strong correlation between values of the NAO for consecutive days, we will focus on autoregressive models. Such models have previously been used by Stephenson et al. (2000) to model the interannual variation of the wintertime seasonal means of the NAO. Stephenson et al. (2000) considered four different stochastic models and compared these to the actual time series by means of different measures. The best fit was found with an autoregressive model taking the means of the 10 last winters into account (the integer 10 was apparently chosen ad hoc), suggesting significant memory effects. We will evaluate the proposed autoregressive models for the NAO by investigating the reliability of forecasts made from the models.

The paper is organized as follows. Section 2 describes the NAO data used and explores basic statistics of the time series. Section 3 discusses the strategy for

autoregressive model fitting of the station-based NAO time series along with a scaling analysis, whereas section 4 discusses the modeling strategy of the noise and implications for forecasting. The dependence of the previous results on the season is discussed in section 5. In section 6, a similar investigation is carried out on the EOF-based time series. Conclusions and a discussion of the results are provided in the last section.

2. Daily indices of the NAO

In this article, the main focus lies on a daily time series of the North Atlantic Oscillation index published by Cropper et al. (2015). The index is calculated from actual sea level pressure (SLP) observations on Iceland and the Azores, but reanalysis data have been used to fill in the gaps (1888–1905, 1940–41, and 145 occasional days) in the observations. This station-based data is freely available,³ and we have chosen to include values between 1 January 1872 and 31 December 2014 in the present study, resulting in 52 230 data points. A visualization of the time series can be found in Fig. 1 (bottom). This figure shows a 3-yr running mean of the observed daily NAO time series. It shows in particular, a clear signature of low-frequency (i.e., decadal) climate variability. In the generation of the time series, the seasonality was removed by applying a tension spline method where a daily annual cycle (of mean SLP and the SLP standard deviation) was interpolated from monthly values and forced so that the average of the daily values of the curve for each month was equal to the monthly means. For details regarding the generation of the time series, we refer to the original source (Cropper et al. 2015).

To emphasize the generality of our results, we have also investigated the time series of daily indices for the NAO provided by the NOAA.⁴ This index is derived by applying rotated principal component analysis (Barnston and Livezey 1987) to monthly standardized 500-hPa geopotential height anomalies obtained from the CDAS in the analysis region 20°–90°N. We have chosen to include values between 1 January 1950 and 31 December 2014 in the present study, resulting in 23 741 data points. The analysis of this time series is deferred to section 6.

We first study the basic properties of the station-based NAO index. We let $\{x_i\}_{i=1}^n$ denote the time series of observed daily NAO index and let n denote the length of the time series, (i.e., $n = 52\,230$). The top-left panel in

² See <http://www.cpc.ncep.noaa.gov/products/predictions/814day/analog.php>.

³ See https://zenodo.org/record/9979#V_9RG037X4i.

⁴ See <http://www.cpc.ncep.noaa.gov/products/precip/CWlink/pna/nao.shtml>.

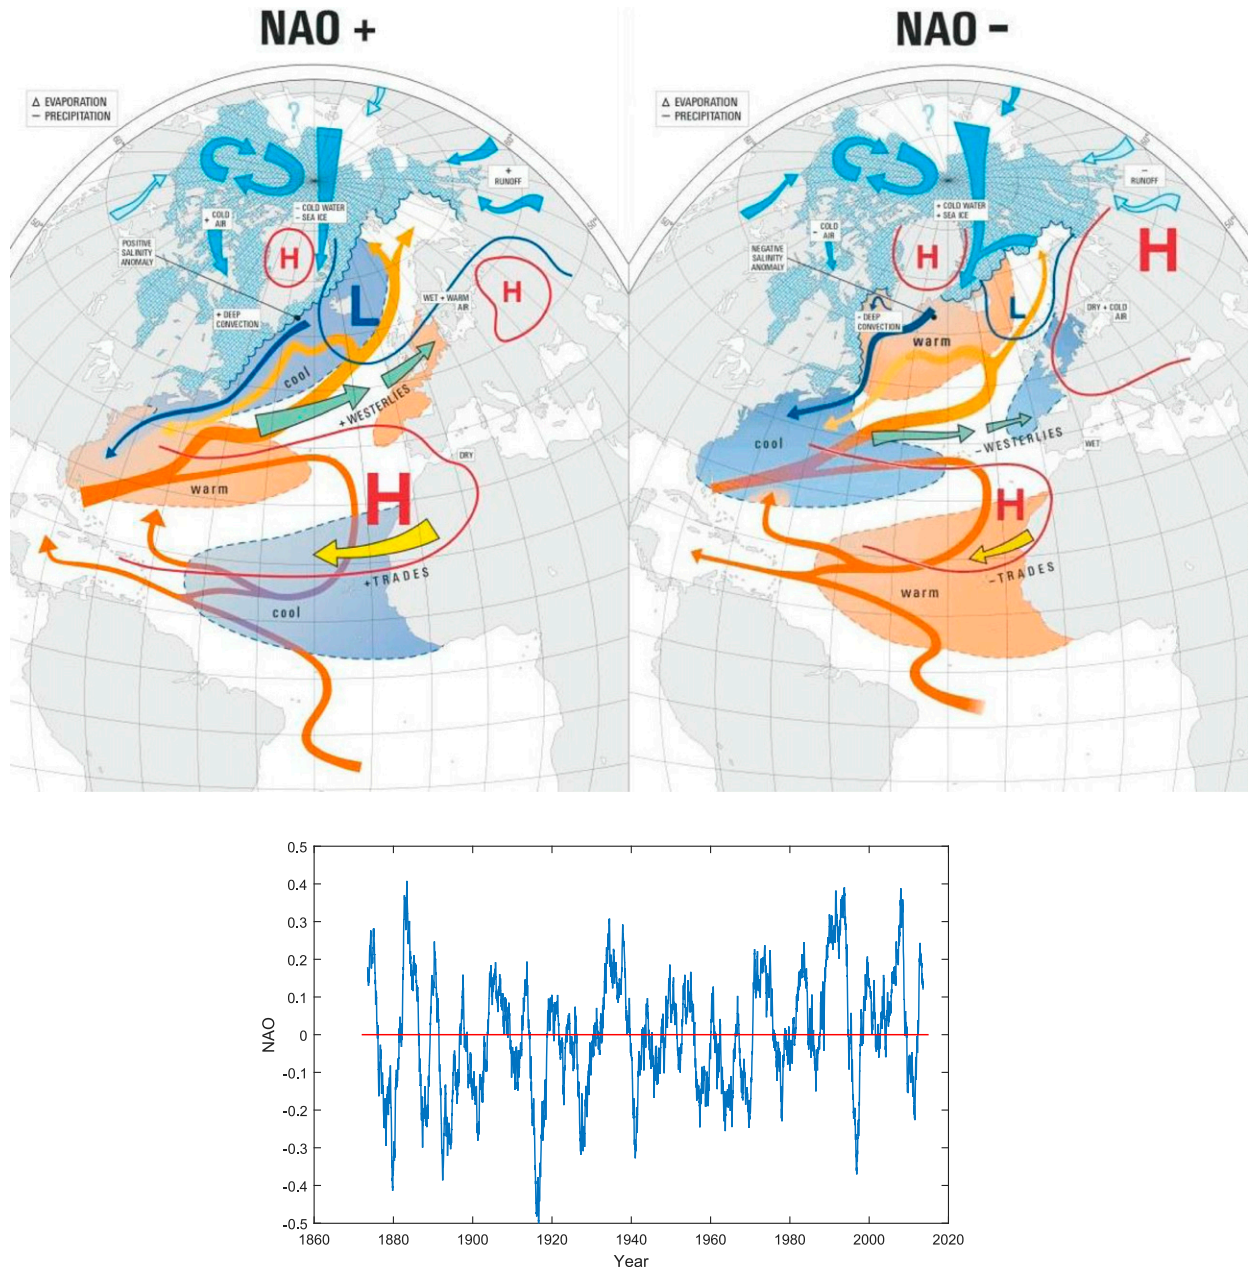


FIG. 1. (top) Adapted from [Wanner et al. \(2001\)](#), shown are typical weather conditions under the positive and negative phases of the NAO. The positive phase of the NAO is characterized by a strong pressure gradient between Iceland and the Azores, yielding stronger westerlies and, in turn, leading to wetter and milder conditions over western and northern Europe and drier conditions over the Mediterranean region. The opposite phase is associated with a weaker pressure gradient and weaker westerlies following a more southward trajectory, resulting in wetter conditions over the Mediterranean region and harsh winters over northern Europe. (bottom) A 3-yr running mean of the observed daily NAO index during the years 1872–2014.

[Fig. 2](#) compares the distribution of the NAO index to a normal distribution with the same mean and variance (-5.9×10^{-5} and 1.614, respectively). The top-right panel in [Fig. 2](#) displays a quantile–quantile (Q–Q) plot of the NAO index. The index distribution shows a couple of characteristic departures from the best-fitted normal distribution. Large positive values and small negative

values are slightly less common, whereas large negative values and positive values on the order of one standard deviation are slightly more common for the NAO data as compared to the normal distribution.

We next turn to the autocovariance of the daily NAO data. The bottom panels in [Fig. 2](#) show the sample autocorrelation function (ACF), which is defined as

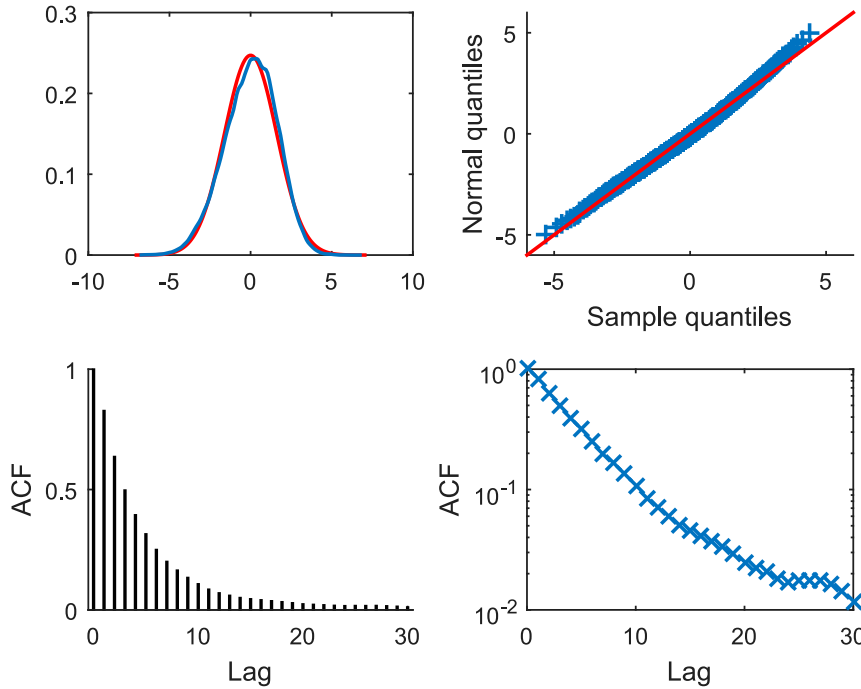


FIG. 2. (top left) The probability density function of the daily NAO index (blue) and normal distributions with the same mean and variance (red). (top right) A Q–Q plot of the daily NAO index compared to a normal sample. (top left) The sample ACF for the daily NAO index in (bottom left) standard scale and (bottom right) logarithmic scale. The 95% confidence bounds for the ACF are on the order of $1.96n^{-1/2} = 8.6 \times 10^{-3}$.

$$\hat{\rho}(h) = \frac{\hat{\gamma}(h)}{\hat{\gamma}(0)}, \quad \text{where} \quad \hat{\gamma}(h) = \frac{1}{n-h} \sum_{i=1}^{n-h} (x_{i+h} - \bar{x})(x_i - \bar{x}),$$

for different values of the time lag h . Here \bar{x} denotes the sample mean of the observations. Clearly, the indices of consecutive days are strongly correlated. From the bottom-right panel in Fig. 2, we also note two different orders of decay of the sample ACF. For lags up to 10 days the decay of the sample ACF is significantly faster than for lags on the order of 10 days or more. We will return to this feature when discussing the mathematical model in section 3b below.

3. Autoregressive models

The form of the sample ACF in Fig. 2 indicates that the daily NAO data could be successfully modeled as an autoregressive (AR) process. An autoregressive process $\{X_i\}_{i=1}^n$ of order p , abbreviated $\text{AR}(p)$, is defined as

$$X_i = \mu + \varphi_1(X_{i-1} - \mu) + \dots + \varphi_p(X_{i-p} - \mu) + \varepsilon_i, \quad \text{for } i \in \{p+1, \dots, n\},$$

where p is a positive integer, $\varphi_1, \dots, \varphi_p$ are real parameters, and μ is the expectation of $\{X_i\}$. Here

$\{\varepsilon_i\}_{i=p+1}^n$ denotes white noise, that is, a sequence of uncorrelated random variables with mean zero and standard deviation σ . There are several ways to determine the optimal choice of order p , and we will investigate a few such methods in the following sections. When p has been chosen, the parameters $\varphi = (\varphi_1, \dots, \varphi_p)$ of the model can be estimated by solving the Yule–Walker equations, which is a linear system of equations given by

$$\hat{\Gamma}_p(\hat{\varphi}_1, \dots, \hat{\varphi}_p)' = \hat{\gamma}_p,$$

where $\hat{\gamma}_p = [\hat{\gamma}_p(1), \dots, \hat{\gamma}_p(p)]'$ and $\hat{\Gamma}_p$ is the matrix $\hat{\Gamma}_p = [\hat{\gamma}_p(i-j)]_{i,j=1}^p$ (Brockwell and Davis 1991). Asymptotically, the parameters $\hat{\varphi} = \hat{\varphi}_1, \dots, \hat{\varphi}_p$ obtained from the Yule–Walker equations approximately satisfy $\hat{\varphi} \in N(\varphi, \sigma^2 \hat{\Gamma}_p^{-1}/n)$, implying that for the NAO time series the error in the parameters is on the order of 10^{-4} , which is insignificant.

a. Partial autocorrelation function

A standard way to determine the optimal value of p for an $\text{AR}(p)$ process designed to model a given time series is by means of the sample partial autocorrelation function (PACF), which is defined as




FIG. 3. Sample PACF as a function of p for the daily NAO index. The size of the statistical error is indicated by the red lines.

$$\hat{\alpha}(h) = \begin{cases} 1, & \text{for } h = 0, \\ \hat{\phi}_{hh}, & \text{for } h \geq 1, \end{cases}$$

where $\hat{\phi}_{hh}$ is the last component of the vector $\hat{\Gamma}_h^{-1}\hat{\gamma}_h$. It can be shown that $\hat{\alpha}(h)$ is the correlation between $x_{i+h} = P_{i,h}(x_{i+h})$ and $x_i = P_{i,h}(x_i)$, with $P_{i,h}(x)$ denoting the projection of x onto the space spanned by $x_{i+1}, \dots, x_{i+h-1}$ (Brockwell and Davis 1991). This means that the sample PACF with lag h can be interpreted as the autocorrelation between x_{i+h} and x_i that cannot be accounted for by autocorrelations with lags $1, \dots, h-1$. For AR(p) processes, the sample PACF at lags greater than p is approximately $N(0, 1/n)$ distributed, so it is natural to choose the integer p as the largest lag for which the modulus of the sample PACF exceeds the statistical error, which should approximately fall within the bounds $\pm 1.96/n^{1/2}$. Figure 3 shows the sample PACF and its statistical error as a function of p . The modulus of the sample PACF is strictly decreasing with p , and it is clear that the optimal choice of p based on the PACF is $p = 3$.

b. Parameter values

Solving the Yule–Walker equations for an AR(3) process using the sample ACF of the daily NAO index as input, the estimated values of the parameters are

$$\hat{\phi} = (0.9626, -0.2042, 0.0534),$$

where the error, as mentioned above, is on the order of 10^{-4} . Figure 4 shows a comparison of the sample ACF of the daily NAO index and of an AR(3) process with the above parameter values. As seen, the AR(3) process provides a very close fit for lags up to 10 days, but for lags on the order of weeks or more, the sample ACF of the NAO is larger than for the model. For lags on the order




FIG. 4. Sample ACF of the daily NAO index (blue) compared to the sample ACF of an AR(3) process forced with nonstandardized t -distributed noise. For the AR(3) process, the median (red \times symbols) and 5%–95% range (red dashed line) of 1000 realizations is displayed.

of months, the difference is up to one order of magnitude. This indicates that there is intraseasonal variability or variability on even longer time scales in the NAO that the AR(3) process is unable to capture. This “shoulder feature” of the sample ACF of the NAO is well known in the literature. Keeley et al. (2009) found that it was mostly due to interannual variability. Franzke and Woollings (2011) found evidence that up to 50% of the winter NAO variability can be explained as arising from climate noise. However, other climate processes might contribute to this shoulder feature and increase the predictive skill of the NAO such as the stratospheric vortex (Scaife et al. 2016) or the Madden–Julian oscillation (Lin et al. 2009). Note, however, as indicated in Fig. 4, that the departure of the sample ACF from that of the AR(3) model is not clearly significant.

A desirable feature of an autoregressive process modeling a time series is that the autoregressive terms explain all dependence in the data and that no long-range dependence can be found in the time series $\{\varepsilon_i\}$ of the remaining noise. In the following two subsections, we will investigate this by two other methods.

c. Aggregated variance

A means to assess if a time series displays long-range dependence is to investigate the sample means of subsets $z_n = \{x_1, \dots, x_n\}$ of the time series for different values of n . If the entries of the time series are independent, it follows from the law of large numbers that the variance of \bar{z}_n is σ^2/n . If the entries are dependent, but only for small time lags, this relation should hold as well, at least for large n . In the presence of long-range




FIG. 5. (left) The sample variance as a function of the sample size for the time series of the daily NAO index (blue) and for the noise in the $AR(p)$ processes modeling the time series (the green line corresponds to $p = 1$ and the black line to $p = 3$). The variances are normalized by division by the sample standard deviation of the entire time series and the reference line (red) is n^{-1} (red). (right) The fluctuation function $F(s)$ as a function of s for the time series of the daily NAO index (blue) and for the noise in the $AR(p)$ processes modeling the time series (the green line corresponds to $p = 1$ and the black line to $p = 3$). The reference line (red) equals $F(s) = s^{1/2}$.

dependence, however, we expect the variance to decay more slowly with n . The left panel in Fig. 5 shows the sample variance of the NAO time series and of the noise in the $AR(3)$ process as a function of sample size. As expected, the short-range dependence of the NAO makes the variance of the sample means considerably higher than σ^2/n . For the $AR(3)$ process, on the other hand, the variance is very close to σ^2/n , but for sample sizes of 20 and larger the variance diverges from σ^2/n indicating long-range dependence in the noise in the $AR(3)$ model.

d. Detrended fluctuation analysis

Long-range dependence of a time series can also be quantified using detrended fluctuation analysis (DFA) (Peng et al. 1994; Taqqu et al. 1995; Bunde et al. 2014; Ludescher et al. 2016). DFA is constructed to detect if the time series is self-affine, in which case it can be rescaled to a self-similar function by a suitable affine transformation, and self-affine time series have dependencies across all time scales. DFA essentially relies on computing

the cumulative sum of the data. To be more precise, we first calculate the cumulative sums

$$z_j = \sum_{i=1}^j (x_i - \bar{x}),$$

where $\{x_i\}$ is the original time series and \bar{x} is its sample mean. The time series $\{z_j\}$ is then divided into n_s non-overlapping segments of lengths s and local least squares straight lines are fit to the cumulative sum on all segments. Subtracting the least squares fits from the cumulative sums, we obtain the detrended series $\{z_j^d\}$. The root-mean-square deviation of $\{z_j^d\}$, known as the fluctuation, is then calculated as

$$F(s) = \left\{ \frac{1}{n_s s} \sum_{k=1}^{n_s} \sum_{j=1}^s [z_{(k-1)s+j}^d]^2 \right\}.$$

The variation of $F(s)$ versus time-scale s reflects the scaling behavior of the time series. If $F(s) \sim s^H$, then H

TABLE 1. Root-mean-square of the difference between the forecasted index m days ahead and the actual value of the index that day.

	1	2	3	4	5	7	10	20
Steady state	0.958	1.389	1.624	1.774	1.884	2.029	2.156	2.239
AR(1)	0.915	1.258	1.411	1.490	1.539	1.587	1.614	1.623
AR(2)	0.905	1.256	1.409	1.489	1.539	1.589	1.616	1.625
AR(3)	0.903	1.253	1.406	1.486	1.536	1.585	1.614	1.624
AR(10)	0.903	1.253	1.406	1.486	1.535	1.585	1.613	1.623

is a generalization of the Hurst exponent. A Hurst exponent in the interval $1/2 < H < 1$ corresponds to a process with positive autocorrelation over large time scales, whereas $H = 1/2$ corresponds to a no-memory process. An AR(p) process modeling a time series explains the long-range dependency of the time series if the Hurst exponent of the remaining noise is close to $1/2$. DFA has been successfully applied to time series originating from a number of different fields, including atmospheric science and hydrology (Kantelhardt et al. 2006; Mandelbrot and Wallis 1968; Maraun et al. 2004; Rust 2007; Hannachi 2014). To give an example, DFA was applied to Northern Ireland daily precipitation data, which was shown to be consistent with a short-memory process (Hannachi 2014). The right panel in Fig. 5 displays the fluctuation $F(s)$ as a function of time-scale s for the NAO index and for the noise in AR(p) processes modeling the index for a number of choices of p . As seen, the original data display clear dependences, whereas the noise in the AR(p) processes for any $p \geq 1$ shows very little long-range dependence. The Hurst exponents for $p \geq 1$ all lie in the interval $[0.49, 0.54]$, implying that there is no obvious self-affinity in the remaining noise. This suggests that the AR(p) process already effectively captures all memory effects that can be detected by means of DFA. But it can also imply that the long-range dependence observed in the sample ACF cannot be captured by DFA, and further analysis may be needed to shed more light on this.

e. Forecasts

In section 3a, we determined the optimal value of p for an AR(p) process modeling the NAO from the PACF, but the optimal p can also be determined by examining forecasts from the models for different values of p . Given values of the NAO index up to time $i - 1$, the best estimate of the index at time i based on an AR(p) model would be

$$x_i = \mu + \theta_1(x_{i-1} - \bar{x}) + \dots + \theta_p(x_{i-p} - \bar{x}),$$

where μ is the sample mean of the time series and $\theta_1, \dots, \theta_p$ are the parameters of the process, which, as mentioned above, can be estimated by the Yule–Walker

equations. Iterating this procedure we can derive forecasts of the index m days into the future. Based on the time series for the NAO, we have calculated the root-mean-square of the difference between the forecasted index m days ahead and the actual value of the index that day. In these forecasts, as well as in the forecasts in section 4c below, we have used the NAO data from the years 1872–1989 as training data to estimate the parameters of the AR(p) process and the NAO data from the years 1990–2014 as testing data to evaluate the forecasts. We note that the difference in parameter estimates between different choices of time interval for the training data is negligible.

It is clear from Table 1 that when p exceeds the optimal $p = 3$ deduced from the PACF, the forecasts do not improve by increasing p . Based only on the performance of the forecasts, we could also settle with $p = 2$ or even $p = 1$. Moreover, the root-mean-square of the predictions based on the AR(p) processes saturate at the sample standard deviation of the testing data ($s_x = 1.621$), and the root-mean-square of the steady state prediction saturates at $2^{1/2}s_x$. Note also that the AR(p) processes are only useful for predicting the NAO index during the coming week and not for dates further into the future.

4. Noise distribution

Figure 2 showed that the distribution of the daily NAO index was similar to a Gaussian distribution but that it possessed certain non-Gaussian features. Unfortunately, one cannot directly apply standard tests of normality to this distribution as such tests assume the data points to be independent. However, applying the Kolmogorov–Smirnov test to a subset of the NAO data containing only every tenth entry (thereby reducing the autocorrelation between two consecutive points to around 10%), we get a p value of around 0.01, so the null hypothesis of normality can be rejected at a low significance level. We will now investigate if the remaining noise in an AR(p) process modeling the daily NAO index is approximately Gaussian or if it more closely resembles some non-Gaussian distribution. The analysis below is based on an AR(3) process, but the results are

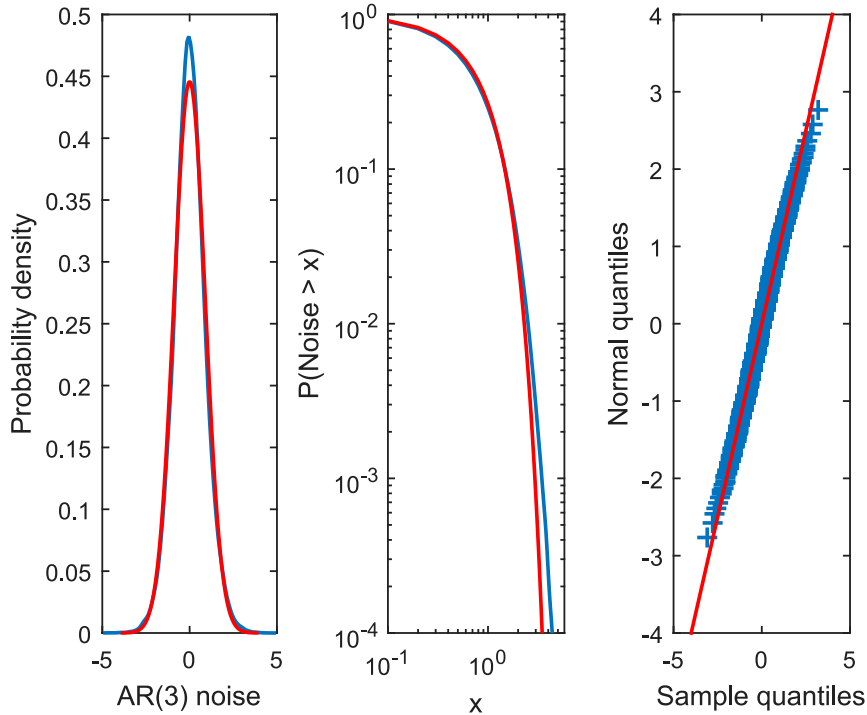


FIG. 6. (left) The probability density function of the noise in the AR(3) model (blue) compared to a Gaussian density (red) with the same mean and variation as the noise. (center) The $P(|X| \geq x)$ as a function of x for the same distributions. (right) A Q–Q plot of the noise distribution compared to a normal sample.

qualitatively the same for all choices of p . Figure 6 compares the distribution of the remaining noise and a normal distribution with the same mean and variance (-9.0×10^{-6} and 0.894, respectively). The noise distribution has a higher frequency of values close to zero and heavier tails than the normal distribution. These features are typical for leptokurtic distributions, that is, distributions with positive excess kurtosis. Recall that the excess kurtosis of a distribution is defined as the kurtosis of the distribution subtracted by the kurtosis of a normal distribution, which is three. The excess kurtosis of the noise distribution is 0.64 as compared to -0.061 for the NAO index. Note also that the difference between the positive and negative tails that was seen in the distribution of the NAO index cannot be found for the noise distribution. In general, the noise distribution is much more symmetric with a skewness of 0.055 compared to -0.245 for the original time series.

a. Leptokurtic distributions

We will next investigate if the noise distribution fits better to some leptokurtic distribution. We consider a broad class of leptokurtic distributions, namely, the α -stable distribution, the nonstandardized t distribution, and the generalized hyperbolic secant distribution. One

could also consider the noncentralized t distribution, but as the noise distribution is very close to being symmetric, the best fits of noncentralized and nonstandardized t distributions are almost indistinguishable. We begin by giving a short description of the leptokurtic distributions considered here.

A random variable X is given by an α -stable distribution with parameters $\alpha \in (0, 2]$, $\beta \in [-1, 1]$, $\gamma \geq 0$, and $\delta \in \mathbb{R}$ if

$$X = \begin{cases} \gamma Z + \delta, & \alpha \neq 1, \\ \gamma Z + \delta + \beta \frac{2}{\pi} \gamma \log \gamma, & \alpha = 1, \end{cases}$$

where Z is a random variable with characteristic function

$$E(e^{iuZ}) = \begin{cases} \exp\left\langle -|u|^\alpha \left[1 - i\beta \tan \frac{\pi\alpha}{2} (\text{sign}u)\right] \right\rangle, & \alpha \neq 1, \\ \exp\left\langle -|u| \left[1 + i\beta \frac{2}{\pi} (\text{sign}u) \log|u|\right] \right\rangle, & \alpha = 1, \end{cases}$$

This form is only one out of several alternative parameterizations of α -stable distributions (Nolan 2016). We note that the Gaussian and Cauchy distributions are

special cases of α -stable distributions with $\alpha = 2$ and $\alpha = 1, \beta = 0$, respectively. Note that, with the exception of $\alpha = 2$, the α -stable distributions are heavy tailed. The decay of the tails is governed by the parameter α as

$$P(|X| > x) \sim x^{-\alpha},$$

for $\alpha \in (0, 2)$ and large x . As a consequence of this tail behavior, the variance and the excess kurtosis are undefined for α -stable distributions with $\alpha \neq 2$. The α -stable distributions play an important role in generalizations of the central limit theorem as the sum of a number of random variables with symmetric distributions having power-law tails (heavy tails), decreasing as $|x|^{-\alpha-1}$, where $\alpha \in (0, 2)$ (and therefore having infinite variance as opposed to the case in the standard central limit theorem), will tend to a stable distribution as the number of summands grows. If the tails of the summands decrease as $|x|^{-k}$ with $k \geq 3$, the sum converges to a stable distribution with $\alpha = 2$ (i.e., a Gaussian distribution). The parameter δ is a location parameter and as the noise distribution for the NAO data is very close to being centered at the origin; we have used $\delta = 0$ throughout the article.

A random variable X is given by a nonstandardized t distribution with parameters $\mu \in \mathbb{R}$, $\sigma > 0$, and $\nu > 0$ if $X = \mu + \sigma T$, where T is a t -distributed random variable with ν degrees of freedom having probability density function

$$f_T(t) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi}\Gamma\left(\frac{\nu}{2}\right)} \left(1 + \frac{t^2}{\nu}\right)^{-\nu+1/2}.$$

The tail decay of this distribution is given by

$$P(|X| > x) \sim x^{-\nu},$$

which is polynomial, but in general faster than for the α -stable distributions as ν can be any positive real number whereas α is bounded from above by 2. The excess kurtosis of the nonstandardized t distribution is finite for $\nu > 4$ and coincides with that of the t distribution and is hence given by $6/(\nu - 4)$, for $\nu > 4$. The nonstandardized t distribution arises in the following way. Assume that the random variable X is normally distributed with mean μ but with unknown, inverse gamma-distributed, variance. Integrating over the gamma distribution, one obtains that X is nonstandardized t distributed. This characterization of the nonstandardized t distribution makes it useful in Bayesian inference, where the inverse gamma distribution is the conjugate prior of the normal distribution with known mean.

TABLE 2. Optimal parameter values for the normal distribution and all three classes of leptokurtic distributions together with their corresponding log likelihoods.

Normal	-68 256
μ	0
σ	0.894
Generalized hyperbolic secant	-68 318
μ	-6.7×10^{-3}
σ	0.936
Nonstandardized t	-67 965
ν	11.9
μ	-3.4×10^{-3}
σ	0.816
α -stable	-68 134
α	1.94
β	0.043
γ	0.610

A random variable X is given by a generalized hyperbolic secant distribution with parameters $\mu \in \mathbb{R}$, and $\sigma > 0$ if $X = \mu + \sigma W$, where W has the probability density function

$$f_W(w) = \frac{1}{2} \operatorname{sech}\left(\frac{\pi w}{2}\right) = \frac{1}{e^{\pi w/2} + e^{-\pi w/2}}.$$

The tail decay of this distribution is given by

$$P(|X| > x) \sim e^{-\pi x/2\sigma},$$

which is exponential, but still not as fast as for the normal distribution, for which $P(|X| > x) \sim e^{-x^2/2\sigma^2}$. The excess kurtosis of the hyperbolic secant distribution is 2. The hyperbolic secant distribution is less common in applications (Ding 2014) but has been included here owing to its tail behavior being somewhere in between the nonstandardized t distributions and the Gaussian distribution.

b. Fit of noise distribution

To determine the best choices of parameters θ within each class of distribution, we will use the maximum likelihood method; that is, we maximize

$$\log L(\theta; \{\varepsilon_i\}_{i=p+1}^n) = \sum_{i=p+1}^n \log f(\varepsilon_i | \theta),$$

where $\{\varepsilon_i\}_{i=p+1}^n$ are the errors in the AR(3) process modeling the daily NAO index. Table 2 shows the optimal parameter values for the normal distribution and all three classes of leptokurtic distributions together with their corresponding log likelihoods. As seen, the

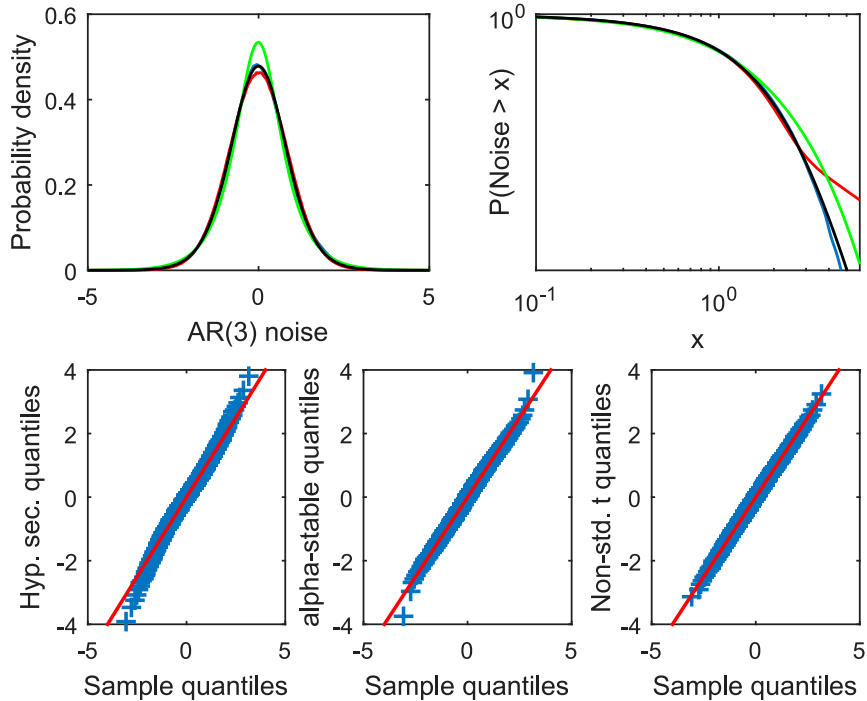


FIG. 7. (top left) The probability density function of the noise of an AR(3) process modeling the daily NAO index (blue) compared to the best-fitted densities of three non-Gaussian distributions, namely, the generalized hyperbolic secant distribution (green), the α -stable distribution (red), and the nonstandardized t distribution (black). (top right) The $P(|X| \geq x)$ as a function of x for the same distributions. (bottom) The Q-Q plots of samples from (left) a generalized hyperbolic secant distribution, (center) an α -stable distribution, and (right) a nonstandardized t distribution vs a sample of the noise in the AR(3) process modeling the daily NAO data.

nonstandardized t distribution gives the largest value of the log-likelihood and hence the best fit.

Figure 7 displays the noise distribution and the best-fitted distributions from each of the three classes of leptokurtic distributions considered and the corresponding Q-Q plots. As seen, the nonstandardized t distribution provides the best overall fit. It is only for values above the 10^{-3} quantile that the distributions differ significantly (the nonstandardized t distribution underestimates the tail slightly). The sums of squares of the residuals in the Q-Q plots with 1000 points are in the interval [1.6, 6.1] for the Gaussian, hyperbolic secant and α -stable distributions but around 0.3 for the nonstandardized t distribution. Moreover, the mean, variance, skewness, and excess kurtosis of the best-fitted nonstandardized t distribution are -3.4×10^{-3} , 0.80, 0, and 0.76, respectively, which is quite close to the sample values stated in the beginning of section 4.

We note, in this context, that the maximum likelihood method can also be used to determine the optimal value of p in the AR(p) models using for example the Aikake information criterion (AIC) or

the Bayesian information criterion (BIC). According to these criteria, one should choose the model that maximizes

$$\text{AIC} = 2k - 2 \log \hat{L} \quad \text{or} \quad \text{BIC} = k \log n - 2 \log \hat{L},$$

respectively, where k is the number of parameters in the model, n is the length of the time series, and \hat{L} is the maximal value of the likelihood function for the model. Comparing $\log \hat{L}$ for noise from a nonstandardized t distribution for different values of p , we observe that for $p = 1, 2,$ and 3 the values of $\log \hat{L}$ are $-68743, -68044,$ and -67965 , respectively, whereas for $p > 3$ it increases slowly with p until it saturates at approximately -67929 for $p > 20$. According to the AIC, the optimal value of p is 21, whereas according to the BIC the optimal value of p is 3, as was deduced also from the PACF. The large disagreement in optimal p between the two criteria reflects the fact that most of the ACF is explained by an AR(3) model (Fig. 4) but that there are correlations on larger time scales as well. As the BIC penalizes the number of parameters more strongly,

TABLE 3. Quantiles x_α of the distribution of the NAO index during the years 1990–2014 for some choices of α .

Probability α	0.001	0.01	0.05	0.1	0.2	0.8	0.9	0.95	0.99	0.999
Quantile x_α	-5.15	-3.89	-2.74	-2.06	-1.28	1.49	2.10	2.61	3.50	4.53

the optimal model deduced by BIC neglects the correlations on larger time scales.

c. Forecast of quantiles

We will next study forecasts of the probability that the NAO index is below a certain quantile m days from now. In contrast to the forecasts of the most likely value in section 3e, the performance of these forecasts will depend on the choice of noise distribution, and as the leptokurtic distributions, in particular the nonstandardized t distribution, provide a better fit of the noise distributions, we expect that the use of leptokurtic noise will improve the quantile forecasts.

We consider a situation where we are interested in the probability that the NAO index switches between a normal value [here quantified as the interval $(\mu - \sigma/2, \mu + \sigma/2) = (-0.81, 0.81)$] to an extreme value m days later. The extreme values are quantified by quantiles in Table 3.

Figures 8 and 9 show the probability that the daily NAO index is below the α quantile m days later for

negative and positive NAO phases, respectively. As seen, the AR(3) process with nonstandardized t -distributed noise is better at predicting the probability of reaching extreme negative states, but in general the difference between different choices of noise distribution is limited. Moreover, the probability that the NAO index predicted by the AR(3) processes lies below a certain quantile saturates at a level that differs significantly from the empirical probabilities. This reflects the fact that the AR(3) process is unable to reproduce the asymmetry of the NAO distribution.

A noteworthy property of the NAO index is that for the positive states (see Fig. 9) the asymptotic probability that the index is below a certain quantile is reached within three to four days, but for the negative states (see Fig. 8), it takes at least 10 days to reach the asymptotic probabilities. This difference is indicative of the different time scales that have been previously observed for the NAO index (Woollings et al. 2010a,b). The fact that the asymptotic probabilities are reached faster for the positive NAO phases is in line with an observed shorter

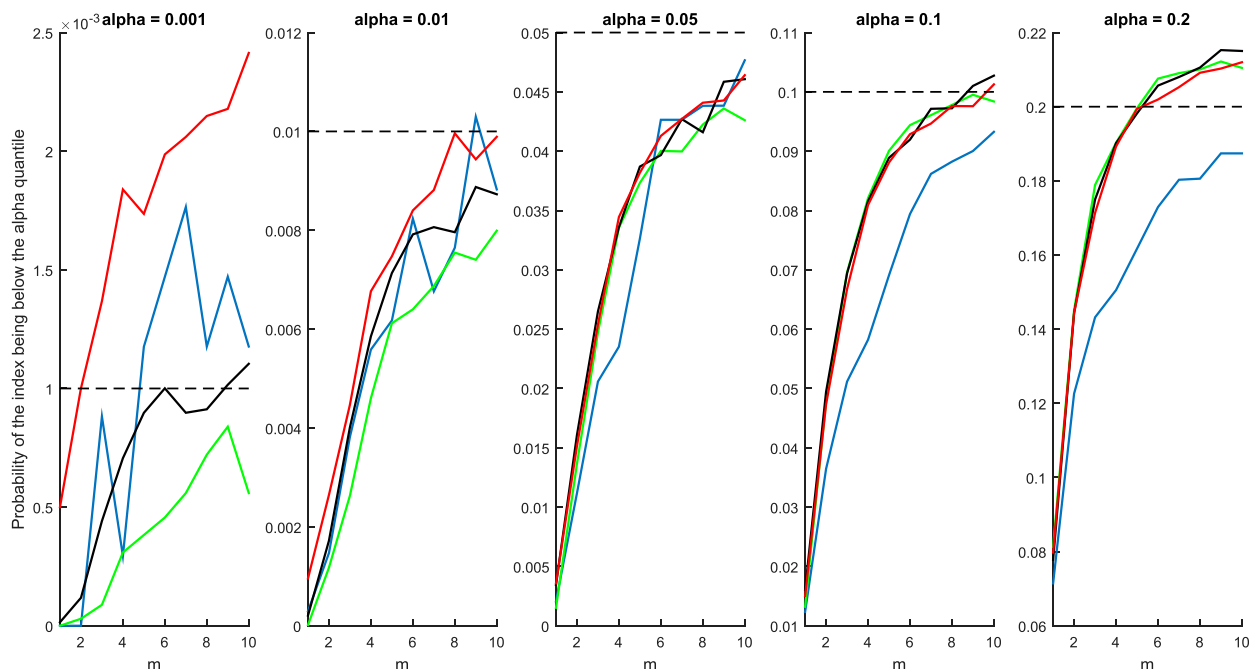


FIG. 8. The plots show the probability that the daily NAO index starting in the interval $[\mu - \sigma/2, \mu + \sigma/2]$ is below the α quantile m days later. Here α take five different values (left)–(right) 0.001, 0.01, 0.05, 0.1, and 0.2, all corresponding to negative NAO indices. The blue lines correspond to the observed index and the other lines to AR(3) processes with Gaussian noise (green), nonstandardized t -distributed noise (black), and α -stable noise (red). The horizontal dashed line represents the asymptotic probabilities.




FIG. 9. As in Fig. 8, but for α values corresponding to positive NAO indices (left)–(right) 0.999, 0.99, 0.95, 0.9, and 0.8.

time scale for the positive NAO phases. The difference between the positive and negative phases of the NAO index cannot be captured by the $AR(p)$ processes as these assume the same autocorrelation for positive and

negative states and the resulting noise distribution is close to being symmetric. In future work, we will address this problem by considering state-dependent and non-linear $AR(p)$ processes.




FIG. 10. The (left) probability density functions and (right) sample ACF of the seasonal daily NAO indices for spring (red), summer (green), autumn (cyan) and winter (purple). The total time series is included for comparison (blue).




FIG. 11. Sample PACF as a function of p for the seasonal daily NAO indices. The size of the statistical error is indicated by the red lines.

5. Seasonal time series

To investigate whether the findings in the preceding sections are independent of the season, the daily NAO data has been divided into four seasonal time series: spring (March–May), summer (June–August), autumn (September–November), and winter (December–February). We have carried out the same analysis as above on these seasonal time series. To give an example of the procedure, the sample ACF of each of the three first seasons has been calculated as the mean of the sample ACFs for all 143 instances of that season found in the data. For the winter data, the sample ACF has been calculated in a similar way, but as there are only 142 full winters in the time series and as the lengths of the winters vary owing to leap years, we have used a weighted mean of the sample ACF for the 142 winters to calculate the winter sample ACF. Figure 10 shows the probability density functions and sample ACFs for the seasonal data. We see that the winter data are more skewed and have longer memory than the yearly data, whereas the other seasons have shorter memory. This is in line with previous observations of persistent NAO states during the winter season (Woollings et al. 2010b; Franzke and Woollings 2011).

Studying the sample PACF of the seasonal time series (Fig. 11) it is clear that $p = 3$ is the optimal choice for all four seasons, just as it was for the total time series. Note,

however, that also in this respect the winter data stand out as for this season the PACF values for lags 2 and 3 are only slightly larger than the statistical error, whereas the values of the PACF for lags 2 and 3 for the other three seasons are considerably larger than the statistical error and also very close to the values of the PACF for the total time series.

We determine the parameters of the AR(3) processes modeling the seasonal data using the Yule–Walker equations. The resulting parameters are found in Table 4. Figure 12 shows Q–Q plots of the noise in the seasonal time series compared to a normal distribution and a nonstandardized t distribution. As can be seen the t distribution provides a better fit. The sums of squares of the residuals in the Q–Q plots with 1000 points are in the interval [1.3, 4.5] for the Gaussian fits, but in the interval [0.11, 0.56] for the nonstandardized t fits. The parameters of the fitted nonstandardized t distributions are given in Table 5.

TABLE 4. Parameter values of the AR(3) processes modeling the seasonal data derived from the Yule–Walker equations.

	θ_1	θ_2	θ_3
Spring	1.008	−0.2186	0.0343
Summer	1.034	−0.3365	0.0890
Autumn	0.9239	−0.1902	0.0625
Winter	0.8834	−0.0688	0.0279
Total	0.9626	−0.2042	0.0534

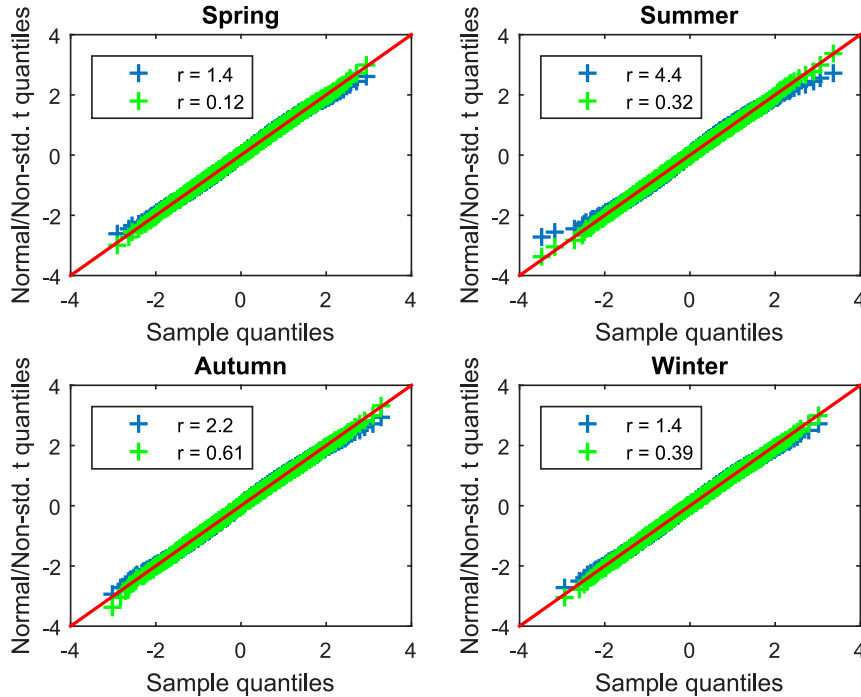


FIG. 12. Q–Q plots of the distribution of the noise in the AR(3) process modeling the seasonal daily NAO indices compared to samples from a normal distribution (blue) and a nonstandardized t distribution (green). The numerical values in the legend of each panel represent the sums of squares of the residuals in the Q–Q plots for the normal distribution (blue) and for the nonstandardized t distribution (green). The red line is the 1:1 line.

6. Analysis of EOF-based time series

We now perform some of the analyses carried out in the preceding sections on the EOF-based NAO index to investigate if there are any qualitative differences between the properties of the two time series. Figure 13 shows the probability density function of the daily NAO index compared to a normal distribution with the same mean and variance. Clearly, the EOF-based NAO index displays the same non-Gaussian features as the station-based NAO index. Considering the sample PACF, its modulus does not decrease as quickly as for the station-based NAO index and the optimal choice of p for a fitted AR process is 5 instead of 3. Fitting an AR(5) process to the EOF-based NAO index, it reproduces the sample ACF very well for lags up to 10 days, but, for lags exceeding two weeks, the sample ACF is significantly larger than for the AR model. We conclude that the long-range dependence in the NAO is more clear for the EOF data than for the station-based data. Finally, Fig. 13 also shows that the resulting noise in the AR(5) model is closely approximated by a nonstandardized t distribution, just as for the station-based data.

7. Summary and conclusions

In this article, we have studied the properties of time series of the daily NAO index, in particular the station-based time series published by Cropper et al. (2015). We have found that the distribution of the NAO index has clear non-Gaussian features and long-range dependence. An autoregressive model with leptokurtic noise taking the values of the index during the last three days into account [AR(3)] provides a good model for the daily NAO index on time scales up to two weeks. Among several leptokurtic distributions considered in this article, the best overall fit is provided by using nonstandardized t -distributed noise in the AR(3) model.

TABLE 5. Optimal parameter values for the nonstandardized t distribution for the noise in the AR(3) processes modeling the seasonal data.

	ν	μ	σ
Spring	14.1	-3.7×10^{-3}	0.792
Summer	8.6	-2.4×10^{-3}	0.779
Autumn	13.1	-2.6×10^{-3}	0.870
Winter	15.1	-6.2×10^{-3}	0.811
Total	11.9	-3.4×10^{-3}	0.816

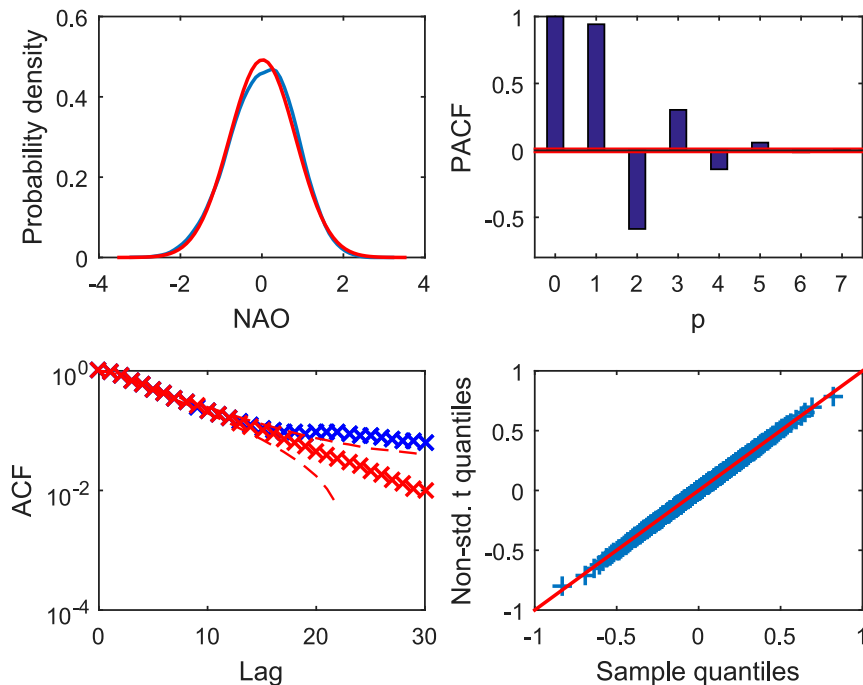


FIG. 13. The plots are based on the EOF based NAO index. (top left) The probability density function of the daily NAO index (blue) and normal distributions with the same mean and variance (red). (top right) The sample PACF as a function of p for the daily NAO index. The size of the statistical error is indicated by the red lines. (bottom left) The sample ACF of the daily NAO index (blue) compared to the sample ACF of an AR(5) process forced with nonstandardized t -distributed noise. For the AR(5) process, the median (red \times symbols) and 5%–95% range (red dashed line) of 1000 realizations is displayed. (bottom right) A Q–Q plot of samples from a nonstandardized t distribution vs a sample of the noise in the AR(5) process modeling the daily NAO data.

Most previous studies of the NAO assumed Gaussian statistics. This assumption is inconsistent with the high nonlinearity and complexity of the processes controlling the climate system. Especially important are the nonlinear interactions between the slow circulation modes, such as the NAO, and fast weather systems. These nonlinear interactions can lead to non-Gaussian behavior and heavy tails (e.g., Majda et al. 2009; Franzke et al. 2005; Franzke 2017; Sura and Hannachi 2015). Here, we show that the NAO has nontrivial non-Gaussian statistics and heavy tails. While systematic procedures predict a multiplicative noise model for the NAO (Majda et al. 2009), here we used a simpler approach using a linear model with additive leptokurtic noise. Our results show that the NAO is a highly nonlinear phenomenon because of its non-Gaussian behavior. However, we were able to develop a linear model with non-Gaussian noise with good predictive skill. Hence, our model can be considered to represent efficiently a weakly multiplicative noise model for which the function of the state variable is weakly dependent on the state variable. Thus, this weakly multiplicative noise is efficiently represented in our

model by a leptokurtic distribution. In future research we will attempt to also derive nonlinear models with explicit multiplicative noise for the NAO and compare their predictive skill with the additive model developed here. Multiplicative noise processes are harder to estimate and require long time series (Peavoy et al. 2015).

Rennert and Wallace (2009) show that the non-Gaussian properties of the NAO might be the results of cross-frequency coupling. They find that the coupling of anomalies with periods >30 days with anomalies with periods of 6–30 days are responsible for most of the skewness of 500-hPa geopotential height. Since skewness is an imprint of non-Gaussianity this mechanism might also cause the here-revealed non-Gaussian characteristics of the NAO.

Luxford and Woollings (2012) provide an alternative explanation for the existence of non-Gaussianity in the atmospheric circulation. They put forward a kinematic argument in which fluctuating jets create deviations from Gaussianity. For instance, geopotential height fields show a distinct pattern of non-Gaussianity: poleward of the jet stream there is positive skewness and equatorward of the

jet stream there is negative skewness. They argue that the non-Gaussian features are a result of a fluctuating jet stream. Since the NAO is intimately linked with the jet streams (Woollings et al. 2010a; Franzke et al. 2011) and the North Atlantic jet stream is an eddy-driven jet stream that is highly variable (e.g., Woollings et al. 2010a; Franzke et al. 2011), this mechanism might also play a role for the non-Gaussian features of the NAO index.

Considering time series for each of the four seasons separately, we have also shown that the results are qualitatively the same as for the entire time series in the sense that an AR(3) model with nonstandardized t -distributed noise provides the best fit to the time series during all four seasons.

We have evaluated the proposed AR(3) model by investigating forecasts of the future distribution of the NAO (considering both the expected value and quantiles) and found that some, but not all, properties are well described by the model. Features that the model is unable to replicate include the long-range dependence on time scales on the order of 20 days or more, the different time scales of the positive and negative phases of the NAO, and the fact that the negative tail of the NAO distributions is fatter than the positive tail. Finding an autoregressive model that includes these features as well is a work in progress.

Our study has focused on the short-time-scale predictive skill of the NAO on the order of a few days. It has recently been shown that the NAO is also predictable on much longer time scales on the order of months up to more than one year (e.g., Domeisen et al. 2015; Dunstone et al. 2016). This predictive skill partly arises from climate modes such as ENSO and the stratospheric vortex. In our future research we will also focus on this long-term behavior of the NAO. Skillful long-term predictions of the NAO are important for many stakeholders.

Acknowledgments. We thank Dr. C. Proistosescu and two anonymous reviewers for their constructive comments that helped to improve the manuscript. CF was supported by the German Science Foundation (DFG) through the Cluster of Excellence “CliSAP” (EXC177), the Collaborative Research Center TRR 181 “Energy Transfers in Atmosphere and Ocean,” and DFG Grant FR3515/3-1.

REFERENCES

- Barnston, A. G., and R. E. Livezey, 1987: Classification, seasonality and persistence of low-frequency atmospheric circulation patterns. *Mon. Wea. Rev.*, **115**, 1083–1126, [https://doi.org/10.1175/1520-0493\(1987\)115<1083:CSAPOL>2.0.CO;2](https://doi.org/10.1175/1520-0493(1987)115<1083:CSAPOL>2.0.CO;2).
- Benedict, J. J., S. Lee, and S. B. Feldstein, 2004: Synoptic view of the North Atlantic oscillation. *J. Atmos. Sci.*, **61**, 121–144, [https://doi.org/10.1175/1520-0469\(2004\)061<0121:SVOTNA>2.0.CO;2](https://doi.org/10.1175/1520-0469(2004)061<0121:SVOTNA>2.0.CO;2).
- Brockwell, P. J., and R. A. Davis, 1991: *Time Series: Theory and Methods*. 2nd ed. Springer, 580 pp.
- Bunde, A., J. Ludescher, C. Franzke, and U. Büntgen, 2014: How significant is West Antarctic warming? *Nat. Geosci.*, **7**, 246–247, <https://doi.org/10.1038/ngeo2126>.
- Colucci, S. J., and D. P. Baumhefner, 1992: Initial weather regimes as predictors of numerical 30-day mean forecast accuracy. *J. Atmos. Sci.*, **49**, 1652–1671, [https://doi.org/10.1175/1520-0469\(1992\)049<1652:IWRAPO>2.0.CO;2](https://doi.org/10.1175/1520-0469(1992)049<1652:IWRAPO>2.0.CO;2).
- Cranz, D., 1765: *Historie von Grönland enthaltend die Beschreibung des Landes und der Einwohner ec. insbesondere die Geschichte der dortigen Mission der Evangelischen Brüder zu Neu Herrnhut und Lichtenfels*. Ebers (Barby), 550 pp.
- Cropper, T., E. Hanna, M. A. Valente, and T. Jónsson, 2015: A daily Azores–Iceland North Atlantic Oscillation index back to 1850. *Geosci. Data J.*, **2**, 12–24, <https://doi.org/10.1002/gdj3.23>.
- Ding, P., 2014: Three occurrences of the hyperbolic-secant distribution. *Amer. Stat.*, **68**, 32–35, <https://doi.org/10.1080/00031305.2013.867902>.
- Domeisen, D., A. Butler, K. Fröhlich, M. Bittner, W. Müller, and J. Baehr, 2015: Seasonal predictability over Europe arising from El Niño and stratospheric variability in the MPI-ESM seasonal prediction system. *J. Climate*, **28**, 256–271, <https://doi.org/10.1175/JCLI-D-14-00207.1>.
- Dunstone, N., D. Smith, A. Scaife, L. Hermanson, R. Eade, N. Robinson, M. Andrews, and J. Knight, 2016: Skillful predictions of the winter North Atlantic Oscillation one year ahead. *Nat. Geosci.*, **9**, 809–814, <https://doi.org/10.1038/ngeo2824>.
- Egede, S. H., 1745: *History of Greenland—A Description of Greenland: Showing the Natural History, Situation, Boundaries, and Face of the Country; the Nature of the Soil; the Rise and Progress of the Old Norwegian Colonies; the Ancient and Modern Inhabitants; Their Genius and Way of Life, and Produce of the Soil; Their Plants, Beasts, Fishes, etc.* (in Danish). Pickering Bookseller, 274 pp.
- Feldstein, S. B., and C. Franzke, 2006: Are the North Atlantic Oscillation and the northern annular mode distinguishable? *J. Atmos. Sci.*, **63**, 2915–2930, <https://doi.org/10.1175/JAS3798.1>.
- , and —, 2017: Atmospheric teleconnection patterns. *Nonlinear and Stochastic Climate Dynamics*, C. L. E. Franzke and T. J. O’Kane, Eds., Cambridge University Press, 54–104.
- Franzke, C., 2009: Multi-scale analysis of teleconnection indices: Climate noise and nonlinear trend analysis. *Nonlinear Processes Geophys.*, **16**, 65–76, <https://doi.org/10.5194/npg-16-65-2009>.
- , 2017: Extremes in dynamic-stochastic systems. *Chaos*, **27**, 012101, <https://doi.org/10.1063/1.4973541>.
- , and T. Woollings, 2011: On the persistence and predictability properties of North Atlantic climate variability. *J. Climate*, **24**, 466–472, <https://doi.org/10.1175/2010JCLI3739.1>.
- , S. Lee, and S. B. Feldstein, 2004: Is the North Atlantic Oscillation a breaking wave? *J. Atmos. Sci.*, **61**, 145–160, [https://doi.org/10.1175/1520-0469\(2004\)061<0145:ITNAOA>2.0.CO;2](https://doi.org/10.1175/1520-0469(2004)061<0145:ITNAOA>2.0.CO;2).
- , A. J. Majda, and E. Vanden-Eijnden, 2005: Low-order stochastic mode reduction for a realistic barotropic model climate. *J. Atmos. Sci.*, **62**, 1772–1745, <https://doi.org/10.1175/JAS3438.1>.
- , T. Woollings, and O. Martius, 2011: Persistent circulation regimes and preferred regime transitions in the North Atlantic. *J. Atmos. Sci.*, **68**, 2809–2825, <https://doi.org/10.1175/JAS-D-11-046.1>.
- , S. Osprey, P. Davini, and N. Watkins, 2015: A dynamical systems explanation of the Hurst effect and atmospheric low-frequency variability. *Sci. Rep.*, **5**, 9068, <https://doi.org/10.1038/srep09068>.
- Graves, T., C. Franzke, N. N. Watkins, R. Gramacy, and E. Tindale, 2017: Systematic Bayesian inference of the long-range dependence and heavy-tail distribution parameters. *Physica A*, **473**, 60–71, <https://doi.org/10.1016/j.physa.2017.01.028>.

- Hannachi, A., 2010: On the origin of planetary-scale extratropical winter circulation regimes. *J. Atmos. Sci.*, **67**, 1382–1401, <https://doi.org/10.1175/2009JAS3296.1>.
- , 2014: Intermittency, autoregression and censoring: A first-order AR model for daily precipitation. *Meteor. Appl.*, **21**, 384–397, <https://doi.org/10.1002/met.1353>.
- , and M. Stendel, 2016: Annex 1: What is NAO? *North Sea Region Climate Change Assessment*, M. Quante and F. Colijn, Eds., Springer, 55–84.
- , T. J. Woollings, and K. Fraedrich, 2012: The North Atlantic jet stream: A look at preferred positions, paths and transitions. *Quart. J. Roy. Meteor. Soc.*, **138**, 862–877, <https://doi.org/10.1002/qj.959>.
- , D. Straus, C. Franzke, S. Corti, and T. Woollings, 2017: Low frequency nonlinearity and regime behavior in the Northern Hemisphere extratropical atmosphere. *Rev. Geophys.*, **55**, 199–234, <https://doi.org/10.1002/2015RG000509>.
- Hurrell, J. W., Y. Kushnir, G. Ottersen, and M. Visbeck, 2003: 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.
- Jung, T., F. Vitart, L. Ferranti, and J. J. Morcrette, 2011: Origin and predictability of the extreme negative NAO winter of 2009/10. *Geophys. Res. Lett.*, **30**, L07701, <https://doi.org/10.1029/2011GL046786>.
- Kantelhardt, J. W., E. Koscielny-Bunde, D. Rybski, P. Braun, A. Bunde, and S. Havlin, 2006: Long-term persistence and multifractality of precipitation river runoff record. *J. Geophys. Res.*, **111**, D01106, <https://doi.org/10.1029/2005JD005881>.
- Keeley, S. P. E., R. T. Sutton, and L. C. Shaffrey, 2009: Does the North Atlantic Oscillation show unusual persistence on intraseasonal timescales? *Geophys. Res. Lett.*, **36**, L22706, <https://doi.org/10.1029/2009GL040367>.
- Lin, H., G. Brunet, and J. Derome, 2009: An observed connection between the North Atlantic Oscillation and the Madden-Julian oscillation. *J. Climate*, **22**, 364–380, <https://doi.org/10.1175/2008JCLI2515.1>.
- Ludescher, J., A. Bunde, C. Franzke, and H. J. Schellnhuber, 2016: Long-term persistence enhances uncertainty about anthropogenic warming of West Antarctica. *Climate Dyn.*, **46**, 263–271, <https://doi.org/10.1007/s00382-015-2582-5>.
- Luxford, F., and T. Woollings, 2012: A simple kinematic source of skewness in atmospheric flow fields. *J. Atmos. Sci.*, **69**, 578–590, <https://doi.org/10.1175/JAS-D-11-089.1>.
- Majda, A. J., C. Franzke, and D. Crommelin, 2009: Normal forms for reduced stochastic climate models. *Proc. Natl. Acad. Sci. USA*, **106**, 3649–3653, <https://doi.org/10.1073/pnas.0900173106>.
- Mandelbrot, B. B., and J. R. Wallis, 1968: Noah, Joseph, and operational hydrology. *Water Resour. Res.*, **4**, 909–918, <https://doi.org/10.1029/WR004i005p00909>.
- Maraun, D., H. W. Rust, and J. Timmer, 2004: Tempting long-memory—On the interpretation of DFA results. *Nonlinear Processes Geophys.*, **11**, 495–503, <https://doi.org/10.5194/npg-11-495-2004>.
- Moore, G. W. K., and I. A. Renfrew, 2012: Cold European winters: Interplay between the NAO and the East Atlantic mode. *Atmos. Sci. Lett.*, **13**, 1–8, <https://doi.org/10.1002/asl.356>.
- Nolan, J. P., 2016: *Stable Distributions—Models for Heavy Tailed Data*. Birkhauser, 352 pp.
- Peavoy, D., C. Franzke, and G. Roberts, 2015: Systematic physics constrained parameter estimation of stochastic differential equations. *Comput. Stat. Data Anal.*, **83**, 182–199, <https://doi.org/10.1016/j.csda.2014.10.011>.
- Peng, C.-K., S. V. Buldyrev, S. Havlin, M. Simons, H. E. Stanley, and A. L. Goldberger, 1994: Mosaic organization of DNA nucleotides. *Phys. Rev.*, **49E**, 1685–1689.
- Pinto, J. G., and C. C. Raible, 2012: Past and recent changes in the North Atlantic oscillation. *Wiley Interdiscip. Rev.: Climate Change*, **3**, 79–90, <https://doi.org/10.1002/wcc.150>.
- Proistosescu, C., A. Rhines, and P. Huybers, 2016: Identification and interpretation of nonnormality in atmospheric time series. *Geophys. Res. Lett.*, **43**, 5425–5434, <https://doi.org/10.1002/2016GL068880>.
- 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, <https://doi.org/10.1175/2009JCLI2669.1>.
- Rust, H. W., 2007: Detection of long-range dependence—Applications in climatology and hydrology. Ph.D. dissertation, Postdam University, 171 pp.
- Sardeshmukh, P. D., and P. Sura, 2009: Reconciling non-Gaussian climate statistics with linear dynamics. *J. Climate*, **22**, 1193–1207, <https://doi.org/10.1175/2008JCLI2358.1>.
- Scaife, A. A., and Coauthors, 2016: Seasonal winter forecasts and the stratosphere. *Atmos. Sci. Lett.*, **17**, 51–56, <https://doi.org/10.1002/asl.598>.
- Stendel, M., E. van den Besselaar, A. Hannachi, E. C. Kent, C. Lefevre, F. Schenk, G. van der Schrier, and T. J. Woollings, 2016: Recent change—Atmosphere. *North Sea Region Climate Change Assessment*, M. Quante and F. Colijn, Eds., Springer, 55–84.
- Stephenson, D. B., V. Pavan, and R. Bojariu, 2000: Is the North Atlantic Oscillation a random walk? *Int. J. Climatol.*, **20**, 1–18, [https://doi.org/10.1002/\(SICI\)1097-0088\(200001\)20:1<1::AID-JOC456>3.0.CO;2-P](https://doi.org/10.1002/(SICI)1097-0088(200001)20:1<1::AID-JOC456>3.0.CO;2-P).
- , H. Wanner, S. Brönnimann, and J. Luterbacher, 2003: The history of scientific research on the North Atlantic Oscillation. *The North Atlantic Oscillation: Climatic Significance and Environmental Impact*, *Geophys. Monogr.*, Vol. 134, Amer. Geophys. Union, 37–50.
- Sura, P., and A. Hannachi, 2015: Perspectives of non-Gaussianity in atmospheric synoptic and low-frequency variability. *J. Climate*, **28**, 5091–5114, <https://doi.org/10.1175/JCLI-D-14-00572.1>.
- Taqqu, M. S., V. Teverovsky, and W. Willinger, 1995: Estimators for long-range dependence: An empirical study. *Fractals*, **3**, 785–798, <https://doi.org/10.1142/S0218348X95000692>.
- Uppala, S. M., and Coauthors, 2005: The ERA-40 Re-Analysis. *Quart. J. Roy. Meteor. Soc.*, **131**, 2961–3012, <https://doi.org/10.1256/qj.04.176>.
- Wallace, J. M., and D. S. Gutzler, 1981: Teleconnections in the geopotential height field during the Northern Hemisphere winter. *Mon. Wea. Rev.*, **109**, 784–812, [https://doi.org/10.1175/1520-0493\(1981\)109<0784:TITGHF>2.0.CO;2](https://doi.org/10.1175/1520-0493(1981)109<0784:TITGHF>2.0.CO;2).
- 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. *Surv. Geophys.*, **22**, 321–382, <https://doi.org/10.1023/A:1014217317898>.
- Woollings, T. J., 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, <https://doi.org/10.1175/2007JAS2347.1>.
- , A. Hannachi, and B. Hoskins, 2010a: Variability of the North Atlantic eddy-driven jet stream. *Quart. J. Roy. Meteor. Soc.*, **136**, 856–868, <https://doi.org/10.1002/qj.625>.
- , —, —, and A. Turner, 2010b: A regime view of the North Atlantic Oscillation and its response anthropogenic forcing. *J. Climate*, **23**, 1291–1307, <https://doi.org/10.1175/2009JCLI3087.1>.
- , C. Franzke, D. Hodson, B. Dong, E. Barnes, C. Raible, and J. Pinto, 2015: Contrasting interannual and multidecadal NAO variability. *Climate Dyn.*, **45**, 539–556, <https://doi.org/10.1007/s00382-014-2237-y>.