University of Groningen

On the predictability of extremes: does the butterfly effect ever decrease?

Sterk, A. E.; Stephenson, D. B.; Holland, M. P.; Mylne, K. R.

Published in:
Quarterly Journal of the Royal Meteorological Society

DOI:
[10.1002/qj.2627](https://doi.org/10.1002/qj.2627)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2016

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Sterk, A. E., Stephenson, D. B., Holland, M. P., & Mylne, K. R. (2016). On the predictability of extremes: does the butterfly effect ever decrease? *Quarterly Journal of the Royal Meteorological Society*, 142(694), 58-64. <https://doi.org/10.1002/qj.2627>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

On the predictability of extremes: Does the butterfly effect ever decrease?

A. E. Sterk,^{a*} D. B. Stephenson,^b M. P. Holland^b and K. R. Mylne^{c†}

^a*Johann Bernoulli Institute, University of Groningen, The Netherlands*

^b*College of Engineering, Mathematics and Physical Sciences, University of Exeter, UK*

^c*Met Office, Exeter, UK*

*Correspondence to: A. E. Sterk, Johann Bernoulli Institute for Mathematics and Computer Science, University of Groningen, PO Box 407, 9700 AK Groningen, The Netherlands. E-mail: a.e.sterk@rug.nl

†The contribution of K. R. Mylne was written in the course of his employment at the Met Office, UK, and is published with the permission of the Controller of HMSO and the Queen's Printer for Scotland.

This study investigates whether or not predictability always decreases for more extreme events. Predictability is measured by the Mean Squared Error (MSE), estimated here from the difference of pairs of ensemble forecasts, conditioned on one of the forecast variables (the 'pseudo-observation') exceeding a threshold.

Using an exchangeable linear regression model for pairs of forecast variables, we show that the MSE can be decomposed into the sum of three terms: a threshold-independent constant, a mean term that always increases with threshold, and a variance term that can either increase, decrease, or stay constant with threshold. Using the generalised Pareto distribution to model wind speed excesses over a threshold, we show that MSE always increases with threshold at sufficiently high threshold. However, MSE can be a decreasing function of threshold at lower thresholds but only if the forecasts have finite upper bounds.

The methods are illustrated by application to daily wind speed forecasts for London made using the 24 member Met Office Global and Regional Ensemble Prediction System from 1 January 2009 to 31 May 2011. For this example, the mean term increases faster than the variance term decreases with increasing threshold, and so predictability decreases for more extreme events.

Key Words: extreme events; predictability; mean squared error; ensemble forecasts

Received 13 December 2014; Revised 21 May 2015; Accepted 1 July 2015; Published online in Wiley Online Library 8 September 2015

1. Introduction

A classical issue in weather prediction is the inherent inaccuracy of forecasts (Tribbia and Anthes, 1987). In one of the earliest studies, Thompson (1957) showed that the predictability of large-scale atmospheric flow patterns is sensitive to the uncertainty in the initial condition. Since the seminal work of Lorenz (1963a), it is well known that deterministic systems can be very unpredictable: small errors in the initial condition can grow exponentially in time. This phenomenon, now known colloquially as *the butterfly effect*,[†] fundamentally limits the accuracy of long-term weather forecasts and stimulated mathematical research on nonlinear dynamics and chaos theory.

Predictability is often quantified in terms of the growth rate of errors in the initial condition. The earliest studies (Smagorinsky,

1963; Mintz, 1964; Leith, 1965) computed the time needed for small errors in the initial condition to double in size. This so-called doubling time is inversely proportional to the largest Lyapunov exponent provided that the initial error is sufficiently small (Lorenz, 2006). However, Lyapunov exponents are independent of the initial condition, which makes them unsuitable as measures for the predictability of specific events such as weather extremes. Finite-time Lyapunov exponents measure the exponential growth rate of nearby trajectories over a finite time and can strongly depend on the initial condition (Nese, 1989; Abarbanel *et al.*, 1991; Smith *et al.*, 1999). Hence, they can be used as a measure of predictability which depends on the event being predicted such as extreme wind speeds at a specific geographical location (Sterk *et al.*, 2012).

Exponential growth rates, as measured by Lyapunov exponents, apply to only small errors and short lead times. Harle *et al.* (2006) showed that for longer lead times the error growth follows a power law which systematically depends on the initial size of the error. For even longer lead times the error will saturate. Another way to quantify predictability is to compute the growth of Mean Squared Error (MSE) between pairs of forecasts starting at different initial conditions (Stephenson and Doblas-Reyes,

[†]or *the seagull effect* when first published: 'When the instability of a uniform flow with respect to infinitesimal perturbations was first suggested as an explanation for the presence of cyclones and anticyclones in the atmosphere, the idea was not universally accepted. One meteorologist remarked that if the theory was correct, one flap of a seagull's wings would be enough to alter the course of the weather forever' (Lorenz, 1963b).

Table 1. Five locations in the United Kingdom for which we have wind speed forecasts.

ID	Location
00003134	Bishopton (Glasgow)
00003355	Church Fenton (Leeds)
00003535	Coleshill (Birmingham)
00003628	Filton (Bristol)
00003772	Heathrow Airport (London)

2000). This squared distance approach is particularly useful in ensemble prediction systems which are not equipped with a tangent linear model needed to compute Lyapunov exponents or their finite-time counterparts. Stroe and Royer (1993) compared different growth formulae for the error between twin forecasts starting from nearby initial conditions. In particular, they showed that the growth rate for small errors and saturation rate for large errors have different characteristics.

Despite much research on forecast error growth, there is hardly any published research on how forecast errors behave for more extreme events. It is often assumed that more extreme events are harder to predict, but this belief needs to be rigorously tested. This study makes an initial attempt to address this problem by investigating how predictability depends on the choice of exceedance threshold used to define an event. We propose simple statistical models for representing ensemble forecasts and use them to find out how the MSE of pairs of forecasts depends upon choice of threshold. The methods are applied to daily wind speed forecasts for London made using the 24 member Met Office Global and Regional Ensemble Prediction System (MOGREPS) from 1 January 2009 to 31 May 2011. Section 2 shows how the MSE depends on lead time and wind speed thresholds. Section 3 then uses extreme value theory (the generalised Pareto distribution, GPD) to quantify the behaviour in the limit of very high thresholds. Section 4 concludes the article with a discussion of our results.

2. Mean squared error growth

2.1. Ensemble forecast data: exchangeable variables

We illustrate our methods using a dataset of ensemble 10 m wind speed forecasts produced by MOGREPS (Bowler *et al.*, 2008). In total, we have 866 forecasts at five locations in the United Kingdom (Table 1) for the period 1 January 2009 until 31 May 2011. For each day, the $m = 24$ ensemble members are initialised at 0600 UTC, and predictions are made out to lead times $\tau = 3, 6, \dots, 54$ h ahead. For brevity, only results for London will be presented here but these are robust: similar results were found at the other four locations.

Let $X_i = X_i(\tau)$ denote the i th ensemble member ($i = 1, \dots, m = 24$) of the wind speed forecasts for a fixed lead time τ . Figure 1 shows box plots of the ensemble members X_i for lead times $\tau = 3$ and $\tau = 54$ h. The box plots suggest that the X_i have similar distributions. It is therefore reasonable to make the assumption that the ensemble members

X_i are exchangeable, which means that $(X_1, \dots, X_m)$ and $(X_{\sigma(1)}, \dots, X_{\sigma(m)})$ have the same joint distribution for any permutation σ of the indices $\{1, \dots, m\}$ (Kallenberg, 2005).

2.2. Relationship between forecasts

The scatter plots in Figure 2 suggest a linear association between X_i and X_j , which can be represented using a linear regression model:

$$X_i = \mu + \rho(X_j - \mu) + \varepsilon_{ij}, \quad (1)$$

where μ and ρ are respectively the mean and the covariance of the X_i s. The error ε_{ij} is considered to be zero mean and constant variance noise which is independent of X_j . For notational simplicity, the dependence of μ , ρ , and ε_{ij} on τ is suppressed. However, we assume that μ and ρ and the variance of the noise are identical for all pairs (i, j) . Point estimates of the μ and ρ parameters can be found easily by using an exchangeable extension of ordinary least-squares estimation (Appendix A).

Figure 3 shows graphs of μ and ρ versus the lead time τ . The mean μ follows a diurnal cycle and attains its minimum at lead times $\tau = 21$ and $\tau = 45$ h which corresponds to 0300 UTC since all forecasts are initialised at 0600 UTC. This agrees with the fact that wind speeds are generally lower during the night. The correlation coefficient ρ (and predictability) decreases with lead time τ , related to the growth of forecast errors.

2.3. Conditional mean squared error

This predictability study adopts the perfect model scenario: instead of comparing forecasts with observations, we will pairwise compare ensemble members with each other. Inclusion of observations would invalidate our assumption of exchangeability because of epistemic uncertainties in how the forecast model represents the real world (e.g. presence of model errors, biases, etc.). One member X_i is assumed to be the ‘forecast’ and another member X_j (with $j \neq i$) is assumed to be a ‘pseudo-observation’. As our measure of predictability, we consider the conditional Mean Squared Error averaged over all distinct pairs of ensemble members:

$$\text{MSE}(\tau, u) = \frac{1}{m(m-1)} \sum_{i \neq j} E([X_i - X_j]^2 | X_j > u), \quad (2)$$

where the threshold u is taken here to be an empirical quantile of all the X_j s.

A key question is how the MSE depends on both the lead time τ and the threshold u . Figure 4(a) shows a graph of the conditional MSE as a function of the lead time τ for fixed thresholds u . For $u = 0$ the MSE is an increasing function of τ . However, for $u > 0$ the MSE has local maxima near $\tau = 21$ and $\tau = 45$ h (0300 UTC) and a local minimum near $\tau = 33$ h (1500 UTC). Thus, errors are larger during the night than during the day. Figure 4(b) shows a graph of the conditional MSE as a function of the threshold u for fixed lead times τ . The MSE increases with u , but the growth rates are larger for $\tau = 21$ and $\tau = 45$ h than for $\tau = 33$ h.

Figure 1. Box plots of the wind speed forecasts for London for lead times (a) $\tau = 3$ h and (b) $\tau = 54$ h.

Figure 2. Scatter plots of the pairs (X_i, X_j) for $1 \leq i, j \leq 3$ for the lead times (a) $\tau = 3$ h and (b) $\tau = 54$ h for London. Scatter plots for the remaining pairs (i, j) are qualitatively similar (not shown).

Figure 3. The coefficients (a) μ and (b) ρ in Eq. (1) for London estimated by the least-squares method described in Appendix A and plotted as a function of the lead time τ . Note that the mean μ follows a diurnal cycle and attains its minimum for $\tau = 21$ h and $\tau = 45$ h (0300 UTC).

But can we understand the reasons for this behaviour in MSE? From Eq. (1) it is shown in Appendix B that the conditional MSE between a pair of forecasts, $E([X_i - X_j]^2 | X_j > u)$, is the sum of three terms:

$$(1 - \rho^2)\text{Var}(X_j) + (1 - \rho)^2 E(X_j - \mu | X_j > u)^2 + (1 - \rho)^2 \text{Var}(X_j | X_j > u). \quad (3)$$

All terms contain the factor $1 - \rho$ or $1 - \rho^2$ and so vanish for perfectly dependent forecasts with $\rho = 1$ (e.g. at initial time $\tau = 0$). The first term involves the unconditional variance of the forecast and so is independent of threshold u . The second term is based on the squared mean of the difference of the forecasts from the mean and so increases monotonically for large threshold $u > \mu$. We will refer to this term as the squared mean excess. The third term is based on the conditional variance of the forecasts, which can either increase or decrease with increasing threshold. If X_j has a finite upper end point, then this term will tend to zero as threshold u approaches the end point.

Ensemble summaries can be obtained by averaging each term over all distinct pairs of ensemble forecasts. For example, for the first term in Eq. (3) we obtain

$$\frac{1}{m(m-1)} \sum_{i \neq j} (1 - \rho^2) \text{Var}(X_j) = \frac{1}{m} \sum_{j=1}^m (1 - \rho^2) \text{Var}(X_j).$$

Analogous expressions hold for the second and third term. To test the validity of the decomposition, Figure 5 shows the difference between the mean square error (Eq. (2)) and the sum of the ensemble averages of the three terms for threshold $u = 0$. Because of sampling variations in the estimates of the components, the difference is not identically zero, however it is only a small fraction

of the MSE and its variation with lead time and threshold shown in Figure 4.

Figure 6 shows the averaged conditional variance and squared excess as functions of lead time τ and threshold u . The conditional variance term increases with τ for fixed u , but decreases with u for fixed τ . The squared mean excess term has local maxima near $\tau = 21$ h and $\tau = 45$ h and a local minimum near $\tau = 33$ h. For fixed lead time, it is an increasing function of u as expected. The mean excess term increases faster than the variance term decreases and so the conditional MSE increases for higher thresholds. In other words, predictability decreases for extremes in this example because of extreme forecasts being further from the mean.

3. What happens at high thresholds?

The decomposition (Eq. (3)) shows that threshold dependence is related to how the first and second moments of forecast excesses vary with threshold (note that for fixed τ the coefficient ρ is a threshold-independent constant). This section will model wind speed excesses over a threshold using extreme value theory to understand the possible behaviour at high thresholds.

3.1. Modelling excesses using the generalised Pareto distribution

Under widely applicable conditions (Balkema and de Haan, 1974; Pickands, 1975) the distribution of excesses $X - u$ is well-approximated by the GPD:

$$G_{\xi, \beta(u)}(x) := \begin{cases} 1 - \left(1 + \frac{\xi x}{\beta(u)}\right)^{-1/\xi} & \text{if } \xi \neq 0, \\ 1 - \exp\left(-\frac{x}{\beta(u)}\right) & \text{if } \xi = 0, \end{cases} \quad (4)$$

Figure 4. (a) Averaged conditional mean square error as a function of the lead time τ for three thresholds u for London. The thresholds $u > 0$ correspond to the 80th and 90th percentile of the distribution of wind speed forecasts. (b) Averaged conditional mean square error as a function of the threshold u for different lead times τ .

Figure 5. The difference between the mean squared error (Eq. (2)) and the decomposition terms in the right-hand side of Eq. (3) as a function of the lead time τ for London.

for X exceeding sufficiently high thresholds. The scale parameter is a linear function of threshold, $\beta(u) = \sigma + \xi(u - \lambda)$, which either increases, stays constant, or decreases depending on whether the shape parameter ξ is positive, zero, or negative, respectively (Coles, 2001). For $\xi < 0$, the distribution has a finite upper end point at $x_F = \lambda - \sigma/\xi$, but is unbounded when the shape parameter is non-negative. The distribution of excesses above any high threshold is completely determined by just three parameters: (λ, σ, ξ) . Note that the location parameter for the tail, λ , should not be confused with the central location parameter $\mu = E(X)$. The GPD can now be used to find the components in the MSE decomposition.

3.2. Conditional MSE components

The first and second moments of the GPD (Appendix C) can be used to find expressions for the mean and variance components[‡] of the MSE decomposition:

$$\begin{aligned} E(u) &= E(X - \mu | X > u)^2 \\ &= \{E(X - u | X > u) + u - \mu\}^2 \\ &= \left\{ \frac{\beta + (1 - \xi)(u - \mu)}{1 - \xi} \right\}^2, \end{aligned}$$

and

$$\begin{aligned} V(u) &= \text{Var}(X | X > u) \\ &= \text{Var}(X - u | X > u) \\ &= \frac{\beta^2}{(1 - \xi)^2(1 - 2\xi)}. \end{aligned}$$

These expressions reveal that the mean component $E(u)$ increases monotonically for $u > \mu + \xi(\lambda - \mu) - \sigma$ regardless of the values of the other parameters. For large u , $E(u) \sim u^2$. In contrast, the variance component $V(u)$ either increases, stays constant, or decreases with u depending upon whether $\xi > 0$, $\xi = 0$, or $\xi < 0$, respectively. Hence, for $\xi \geq 0$, the sum of the two terms $S(u) = E(u) + V(u)$ is a strictly monotonically increasing

function of threshold: the butterfly effect increases for higher thresholds. For $\xi < 0$, decreasing $V(u)$ compensates increasing $E(u)$ to create a minimum $S(u)$ at

$$u_{\min} = \frac{\mu(1 - 2\xi) + \xi\lambda - \sigma}{1 - \xi}. \quad (5)$$

As u tends to x_F , $V(u)$ decreases at a slower rate than $E(u)$ increases and so $S(u)$ increases as the threshold approaches the end point. In summary, the MSE always increases with threshold at sufficiently high threshold, although it can be a decreasing function of threshold at lower thresholds, but only if the forecasts have finite upper bounds ($\xi < 0$).

It is instructive to consider some special cases. Figure 7 shows the behaviour of the components when the forecasts are simulated from uniform and normal distributions. The results agree with what is expected from the above GPD theory. The $U(0, 1)$ uniform distribution can easily be shown to have $\lambda = 0$, $\sigma = 1$, $\xi = -1$ and $\mu = 0.5$ and so one expects an increasing mean component, a decreasing variance component and a minimum at $u_{\min} = 0.25$, as can be seen in Figure 7(a). The normal distribution can be shown to have $\xi = 0$ and so one expects to see an increasing mean component and a constant variance component, which is occurring for larger threshold in Figure 7(b).

3.3. GPD analysis of the wind speed data

Can we use GPD fits to understand the behaviour seen earlier in our wind speed forecast example?

The parameter μ can be estimated from the linear model (1) using the least-squares procedure described in Appendix A. The parameters (λ, σ, ξ) can be estimated by first selecting a fixed threshold u_0 and then fitting the exceedances $X - u_0 > 0$ to a Poisson process (Coles, 2001). This procedure is implemented in the `fplot` command, which is part of the R package `evd` (Stephenson, 2002).

The selection of the threshold u_0 depends on two contradicting requirements. On the one hand u_0 should be as small as possible in order to have many exceedances $X - u_0 > 0$ thereby decreasing the variance of the estimates for (λ, σ, ξ) . On the other hand u_0 should be large in order to ensure that the asymptotic GPD is a valid model for the exceedances $X - u_0 > 0$. If $G_{\beta, \xi}$ is a valid model for $X | X > u_0$, then the mean excess function $E(X - u | X > u)$ for $u > u_0$ is linear in u . Moreover, the shape parameter ξ and the modified scale parameter $\beta - \xi u$ should be approximately linear in u . We adopt these criteria for selecting u_0 .

For example, Figure 8 shows threshold selection plots for the variable X_1 of the London wind speed forecasts for the lead time $\tau = 21$ h. Based on these plots we set $u_0 = 3.5$. Subsequently, the parameters (λ, σ, ξ) can be estimated. The diagnostic plots in Figure 9 indicate that the fit is fairly good: the quantile plot is nearly diagonal and the empirical return levels agree well with the fitted return levels. This procedure has been repeated for the lead times $\tau = 33$ and $\tau = 45$ h (Table 2).

The wind speed forecasts have a negative shape parameter which implies that the function $S(u)$ might have a minimum. However, the value $u_{\min}$ given in Table 2, computed from the

[‡]without the ρ -dependent multipliers.

Figure 6. (a) The conditional variance $(\rho - 1)^2 \text{Var}(X_j | X_j > u)$ and (b) squared excess $(\rho - 1)^2 E(X_j - \mu | X_j > u)$ (both averaged over all ensemble members) as a function of the lead time τ for fixed thresholds u . (c,d) show the same quantities as a function of the threshold u for fixed lead times.

Figure 7. Examples of threshold dependence in $S(u) = E(u) + V(u)$ (solid line), $E(u)$ (dashed line), and $V(u)$ (dotted line) for forecasts that are (a) uniformly distributed $X \sim U(0, 1)$, and (b) normally distributed $X \sim N(0, 1)$. Curves were obtained by simulating 10^5 values from the respective distributions.

GPD parameter estimates using Eq. (5), is far below the average of the wind speed forecasts (Figure 1). At such low values, the GPD is no longer a good approximation for excesses, which may help to explain why no such minimum can be seen in Figure 4.

4. Conclusion

This study has investigated whether or not predictability can increase for more extreme events. Predictability is measured

here by the Mean Squared Error (MSE) between pairs of ensemble forecasts, conditioned on one of the forecast variables (the ‘pseudo-observation’) exceeding a threshold. Assuming an exchangeable linear regression model for pairs of forecast variables, we have found that the MSE can be decomposed into the sum of three terms: a threshold-independent constant, a mean term that always increases with threshold, and a variance term that can either increase, decrease, or stay constant with threshold. Extreme value theory has revealed that MSE always increases with

Figure 8. Threshold selection plots for the variable X_1 and lead time $\tau = 21$ h; (a) mean excess, (b) modified scale and (c) shape parameters.

Figure 9. (a) Quantile plot and (b) return level plot obtained from estimating the parameters (λ, σ, ξ) for the lead time $\tau = 21$ h. The parameters were estimated using 866 daily data values and so the return period, which is the reciprocal of the exceedance probability, is in units of 866 days.

Table 2. Point process estimates for the parameters (λ, σ, ξ) for different lead times τ .

τ	u_0	$\hat{\lambda}$	$\hat{\sigma}$	$\hat{\xi}$	$u_{\min}$
21	3.5	9.56 (0.50)	0.69 (0.16)	−0.13 (0.04)	2.00
33	4.8	9.74 (0.25)	0.29 (0.06)	−0.32 (0.04)	2.90
45	3.5	10.17 (0.64)	0.91 (0.22)	−0.07 (0.05)	2.04

The standard errors are listed between parentheses.

threshold at sufficiently high threshold. Furthermore, MSE can be a decreasing function of threshold at lower thresholds but only if the forecasts have finite upper bounds.

In summary, our modelling approach suggests that the butterfly effect generally increases for more extreme events. This is predominantly due to the average distance between pairs of forecasts increasing if one of the forecasts is constrained to be in a more extreme edge of state space. For bounded forecasts, a reduction in variance can help to compensate this but is insufficient to do so as one approaches the edge of the attractor.

The methods have been applied to daily wind speed forecasts for London made using the 24 member MOGREPS from 1 January 2009 to 31 May 2011. The MSE has a local minimum during the nights and a local maximum during the afternoons. The MSE is found to increase monotonically with threshold, especially so during the night. The mean component increases and the variance component decreases with threshold as is to be expected for variables that are estimated to be bounded from above (negative GPD shape parameter).

Our approach relies on several assumptions, which appear reasonable but should be investigated in more detail in future work. For example, we have modelled dependency between forecasts using a simple linear model with constant variance noise. This appears to work well for our data example but more complex models may be appropriate for other datasets. It is possible that there could be more predictability at higher thresholds if one were to allow for extremal dependency between forecasts using bivariate copula approaches (Stephenson *et al.*, 2008).

This study has addressed predictability by focussing on mean squared error – a measure of skill for deterministic forecasts. In future studies, it would be of interest to extend the approach to address probability forecasts e.g. by using modified decompositions to decompose the Brier score, which for frequency of exceedance above a threshold may be expressed as a mean squared error of binary forecasts. One other caveat is that this study has measured predictability using an absolute score but one may find alternative conclusions if one were to assess predictability by measuring skill against no-skill forecasts (e.g. using the MSE skill score against either persistence or climatological forecasts).

As with all previous perfect model approaches, we have made the assumption that observations are exchangeable with ensemble forecasts. When this is not the case and the forecasts contain more unpredictable noise than the observations (e.g. Eade *et al.*, 2014), then real-world observations will give smaller MSE than the MSE of pairs of forecasts. However, we suspect that the broad conclusions about increase of MSE with threshold may still

generally apply to unreliable forecasts that are co-exchangeable with the observations.

Although we have illustrated our approach using weather forecasts of wind speeds, the methodology is generic and so could be applied to other predicted quantities of interest such as temperatures or river flows.

Acknowledgements

This research has been carried out within the project ‘PREDEX: PREdictability of EXtreme weather events’, funded by Complexity-NET (<http://www.complexitynet.eu>). The authors gratefully acknowledge support by the UK and Dutch funding agencies involved in Complexity-NET: the Engineering and Physical Sciences Research Council (EPSRC) and the Netherlands Organisation for Scientific Research (NWO).

The authors wish to thank Dr Renato Vitolo for his contributions to the project, Dr Lisa Murray (Met Office) for preparing the data and Dr Jan Barkmeijer (Royal Netherlands Meteorological Institute) for useful discussions. The authors also wish to thank the three reviewers for their insightful comments which have helped improve the manuscript.

Appendix A

Ordinary least-squares estimation for exchangeable variables

The coefficients μ and ρ in Eq. (1) are estimated by a least-squares method as follows. First, rewrite Eq. (1) as

$$X_i = a + bX_j + \varepsilon_{ij},$$

where $a = (1 - \rho)\mu$ and $b = \rho$. The coefficients a and b are obtained by minimizing the expression

$$\text{SSE} = \frac{1}{N} \sum \{\mu + \rho(X_j - \mu) - X_i\}^2,$$

where the summation ranges over all 866 days and all pairs $i \neq j$ so that $N = 866 \cdot m(m - 1)$. Setting the partial derivatives of the SSE with respect to μ and ρ to 0 gives the equations

$$\begin{aligned} Na + b \sum X_j &= \sum X_i, \\ a \sum X_j + b \sum X_j^2 &= \sum X_i X_j, \end{aligned}$$

from which a and b are readily solved. Finally, $\rho = b$ and $\mu = a/(1 - b)$. This computation is repeated for each lead time τ .

Appendix B

Decomposition of the MSE

Recall the formula for conditional variance (Ross, 1997):

$$\begin{aligned} \text{Var}(X | Y) &= E(\{X - E(X | Y)\}^2 | Y) \\ &= E(X^2 | Y) - E(X | Y)^2. \end{aligned}$$

Using Eq. (1), the MSE of $X_i - X_j$ conditional on $X_j > u$ can be written as

$$\begin{aligned} E([X_i - X_j]^2 | X_j > u) &= \text{Var}(X_i - X_j | X_j > u) + E(X_i - X_j | X_j > u)^2 \\ &= \text{Var}\{(1 - \rho)\mu + (\rho - 1)X_j + \varepsilon_{ij} | X_j > u\} \\ &\quad + E\{(1 - \rho)\mu + (\rho - 1)X_j + \varepsilon_{ij} | X_j > u\}^2 \\ &= \text{Var}\{(\rho - 1)X_j + \varepsilon_{ij} | X_j > u\} \\ &\quad + \{(\rho - 1)E(X_j - \mu | X_j > u) + E(\varepsilon_{ij} | X_j > u)\}^2. \end{aligned}$$

The latter expression simplifies under the additional assumption that the error ε_{ij} is independent of X_j . This implies that

$$E(\varepsilon_{ij} | X_j > u) = E(\varepsilon_{ij}) = 0$$

and

$$\begin{aligned} \text{Var}((\rho - 1)X_j + \varepsilon_{ij} | X_j > u) &= (\rho - 1)^2 \text{Var}(X_j | X_j > u) + \text{Var}(\varepsilon_{ij} | X_j > u) \\ &= (\rho - 1)^2 \text{Var}(X_j | X_j > u) + \text{Var}(\varepsilon_{ij}). \end{aligned}$$

Note that Eq. (1) implies that

$$\text{Var}(X_i) = \rho^2 \text{Var}(X_j) + \text{Var}(\varepsilon_{ij}).$$

Assuming that $\text{Var}(X_i) = \text{Var}(X_j)$ gives

$$\text{Var}(\varepsilon_{ij}) = (1 - \rho^2) \text{Var}(X_i) = (1 - \rho^2) \text{Var}(X_j).$$

Hence, the final decomposition of the conditional mean square error reads as

$$\begin{aligned} E([X_i - X_j]^2 | X_j > u) &= (1 - \rho^2) \text{Var}(X_j) + (\rho - 1)^2 \text{Var}(X_j | X_j > u) \\ &\quad + (\rho - 1)^2 E(X_j - \mu | X_j > u)^2. \end{aligned}$$

Appendix C

Moments of the generalised Pareto distribution

Let X have a GPD with scale parameter β and tail index ξ . For any natural number r satisfying $\xi < 1/r$ we have

$$E(X^r) = \frac{\beta^r}{\xi^{r+1}} \cdot \frac{\Gamma(\xi^{-1} - r)}{\Gamma(\xi^{-1} + 1)} r!,$$

where Γ is the Gamma function (Embrechts *et al.*, 1997). Note that the expectation of X^r does not exist when $1/r \leq \xi$. For $\xi < 1/2$ we have

$$E(X) = \frac{\beta}{1 - \xi} \quad \text{and} \quad E(X^2) = \frac{2\beta^2}{(1 - \xi)(1 - 2\xi)}.$$

Hence, the variance is given by

$$\text{Var}(X) = \frac{\beta^2}{(1 - \xi)^2(1 - 2\xi)}.$$

Furthermore, the mean excess function is given by

$$E(X - u | X > u) = \frac{\beta + \xi u}{1 - \xi}.$$

References

- Abarbanel H, Brown R, Kennel M. 1991. Variation of Lyapunov exponents on a strange attractor. *J. Nonlinear Sci.* **1**: 175–199.
- Balkema A, de Haan L. 1974. Residual life time at great age. *Ann. Probab.* **2**: 792–804.
- Bowler N, Arribas A, Mylne KR, Robertson K, Beare S. 2008. The MOGREPS short-range ensemble prediction system. *Q. J. R. Meteorol. Soc.* **34**: 703–722.
- Coles S. 2001. *An Introduction to Statistical Modeling of Extreme Values*, Springer Series in Statistics. Springer: New York, NY.
- Eade R, Smith D, Scaife A, Wallace E, Dunstone N, Hermanson L, Robinson N. 2014. Do seasonal-to-decadal climate predictions underestimate the predictability of the real world? *Geophys. Res. Lett.* **41**: 5620–5628, doi: 10.1002/2014GL061146.
- Embrechts P, Klüppelberg C, Mikosch T. 1997. *Modelling Extremal Events, Applications of Mathematics: Stochastic Modelling and Applied Probability*, 33. Springer-Verlag: Berlin and Heidelberg, Germany.
- Harle M, Kwasniok F, Feudel U. 2006. Growth of finite errors in ensemble prediction. *Nonlin. Processes Geophys.* **13**: 167–176.
- Kallenberg O. 2005. *Probabilistic Symmetries and Invariance Principles*, Probability and its Applications, Springer-Verlag: Berlin and Heidelberg, Germany.
- Leith CE. 1965. Numerical simulation of the Earth's atmosphere. In *Methods in Computational Physics*, Alder B, Fernbach S, Rotenberg M. (eds) 1–28. Academic Press/Elsevier: Amsterdam, Netherlands.
- Lorenz EN. 1963a. Deterministic nonperiodic flow. *J. Atmos. Sci.* **20**: 130–141.
- Lorenz EN. 1963b. The predictability of hydrodynamic flow. *Trans. NY Acad. Sci.* **25**: 409–432.
- Lorenz EN. 2006. Predictability – A problem partly solved. In *Predictability of Weather and Climate*, Palmer TN, Hagedorn R. (eds.): 40–58. Cambridge University Press: Cambridge, UK.
- Mintz Y. 1964. 'Very long term global integration of the primitive equations of atmospheric motion', WMO-IUGG Symposium on research and development aspects of long-range forecasting, Boulder, Colorado. Technical Note 66. WMO: Geneva, Switzerland, 141–155.
- Nese J. 1989. Quantifying local predictability in phase space. *Physica D* **35**: 237–250.
- Pickands J. 1975. Statistical inference using extreme order statistics. *Ann. Stat.* **3**: 119–131.
- Ross S. 1997. *A First Course in Probability* (5th edn). Prentice Hall: Upper Saddle River, NJ.
- Smagorinsky J. 1963. General circulation experiments with the primitive equations. I. The basic experiment. *Mon. Weather Rev.* **91**: 99–164.
- Smith L, Ziehmann C, Fraedrich K. 1999. Uncertainty dynamics and predictability in chaotic systems. *Q. J. R. Meteorol. Soc.* **125**: 2855–2886.
- Stephenson A. 2002. 'evd: Extreme value distributions'. *R News* **2/2** 31–32. <http://CRAN.R-project.org/doc/Rnews/> (accessed 1 August 2015).
- Stephenson D, Doblas-Reyes F. 2000. Statistical methods for interpreting Monte Carlo ensemble forecasts. *Tellus A* **52**: 300–322.
- Stephenson D, Casati B, Ferro C, Wilson C. 2008. The extreme dependency score: A non-vanishing measure for forecasts of rare events. *Meteorol. Appl.* **15**: 41–50.
- Sterk A, Holland M, Rabassa P, Broer H, Vitolo R. 2012. Predictability of extreme values in geophysical models. *Nonlin. Processes Geophys.* **19**: 529–539.
- Stroe R, Royer J. 1993. Comparison of different error growth formulas and predictability estimation in numerical extended-range forecasts. *Ann. Geophys.* **11**: 296–316.
- Thompson P. 1957. Uncertainty of initial state as a factor in the predictability of large-scale atmospheric flow patterns. *Tellus* **9**: 275–295.
- Tribbia J, Anthes R. 1987. Scientific basis of modern weather prediction. *Science* **31**: 493–499.