

Effects of stochastic parametrization on extreme value statistics

Article

Accepted Version

Hu, G. ORCID: <https://orcid.org/0000-0003-4305-3658>, Bodai, T. and Lucarini, V. ORCID: <https://orcid.org/0000-0001-9392-1471> (2019) Effects of stochastic parametrization on extreme value statistics. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29 (8). 083102. ISSN 1089-7682 doi: 10.1063/1.5095756 Available at <https://centaur.reading.ac.uk/85432/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1063/1.5095756>

Publisher: American Institute of Physics

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

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Effects of Stochastic Parametrization on Extreme Value Statistics

Guannan Hu,^{1,2} Tamás Bódai,³ and Valerio Lucarini^{3,4,2}

¹School of Integrated Climate System Sciences (SICSS), University of Hamburg, Hamburg, Germany

²CEN, University of Hamburg, Hamburg, Germany

³Department of Mathematics and Statistics, University of Reading, Reading, UK

⁴Centre for the Mathematics of Planet Earth, University of Reading, Reading, UK

Abstract

Extreme geophysical events are of crucial relevance to our daily life: they threaten human lives and cause property damage. To assess the risk and reduce losses, we need to model and probabilistically predict these events. Parametrizations are computational tools used in Earth system models, which are aimed at reproducing the impact of unresolved scales on resolved scales. The performance of parametrizations has usually been examined on typical events rather than on extreme events. In this paper we consider a modified version of the two-level Lorenz'96 model and investigate how two parametrizations of the fast degrees of freedom perform in terms of the representation of extreme events. One parametrization is constructed following Wilks (2005) and is constructed through an empirical fitting procedure; the other parametrization is constructed through the statistical mechanical approach proposed by Wouters and Lucarini (2012, 2013). The two strategies show different advantages and disadvantages. We discover that the agreement between parametrized models and true model is in general worse when looking at extremes rather than at the bulk of the statistics. The results suggest that stochastic parametrizations should be accurately and specifically tested against their performance on extreme events, as usual optimization procedures might neglect them.

1 Introduction

Weather and climate models are mathematical representations of the physical processes in the Earth system. These physical processes operate on different temporal and spatial scales. For instance, atmospheric convection acts from minutes to hours on a spatial scale of kilometers, while ocean circulation acts from years to multiple thousands of years on a spatial scale of thousands of kilometers. Since weather and climate models have certain temporal and spatial resolutions, the processes that happen on shorter timescales or smaller spatial scales cannot be resolved by the model. However, the processes on different scales interact with each other, and the unresolved processes will influence the resolved processes. Parametrizations are computational tools aimed at reproducing the impact of the unresolved scales on the scales that can be directly resolved. Different parametrization schemes strike different balance between computational efficiency and various aspects of accuracy. While traditionally parametrizations are formulated as empirical deterministic formulas, in recent years the scientific community has advocated the need for using more general parametrization methods comprising also of stochastic components; see Palmer and Williams (2008), Franzke *et al.* (2015) and Berner *et al.* (2017) for a summary of recent developments in this field. The idea of using a stochastic climate model has been proposed by Hasselmann (1976). The goal of this paper is to address, in a simple yet informative case, the problem of how well parametrizations perform, but, instead of looking at the bulk of the statistics, as usually done, we look at extreme value statistics.

The understanding, modeling and probabilistic prediction of extremes are of key interest to the financial market, the insurance sector, and also to civil defence services concerning natural catastrophes, e.g. hurricanes, storms, floods, etc. A mature statistical framework, known as extreme value theory (EVT), is widely applied to analyze extremes. The theory was developed in the course of the 20th century by Fisher and Tippett (1928), Gnedenko (1943), Balkema and de Haan (1974), and Pickands (1975); an excellent summary of the main results can be found in, e.g. Embrechts,

Klüppelberg, and Mikosch (1997), and Coles (2001). Two approaches are often used to analyze extremes: the block maxima (BM) and the peak over threshold (POT) approaches. Extremes are defined in different ways by these two approaches: the first one takes the maximal values in blocks or batches of sample data as extremes, while the second one selects the exceedances above a given high threshold. Under rather general conditions the block maxima are distributed according to the Generalized Extreme Value (GEV) distribution, and the threshold exceedances follow the Generalized Pareto (GP) distribution. That is in the limit of large block sizes and thresholds, respectively, upon some suitable normalisation. Holland *et al.* (2012) applied the BM method to the extremes of a number of chaotic deterministic dynamical systems and found that for sufficiently smooth so-called physical observables, a parameter of the limiting GEV distribution, the shape parameter, is determined by the dimensions of the stable and unstable manifolds of the chaotic attractor. Later Lucarini *et al.* (2014) applied the POT method and suggested by a heuristic argument that the formula of Holland *et al.* (2012) should apply generically. These two papers together with other earlier publications Collet (2001); Freitas *et al.* (2010) establish a connection between extreme value statistics and the geometrical properties of the attractor. A comprehensive summary of the main results of extreme events of observables of deterministic systems with examples and applications can be found in Lucarini *et al.* (2016). This link has been reexamined by Gálfi, Bódai, and Lucarini (2017), in which the authors presented the convergence of shape parameter estimates to the theoretical value in a two-level quasi-geostrophic atmospheric model, or the lack of it, as this convergence could be observed only in the model with a strong forcing. Furthermore, Bódai (2017) argued that the convergence of the shape parameter can be observed typically for high-dimensional systems, and in low-dimensional systems, such as the Lorenz'84 and one-level Lorenz'96 models that he studied, the shape parameter estimates can increase nonmonotonically with the block size, or fluctuate owing to the fractality of the natural measure, in which latter case no extreme value law exists in a strict sense.

We use a conceptual model of the atmosphere, the two-level Lorenz'96 model (L96) Lorenz (1995) and two parametrized versions of this model, which are constructed by resolving the large-scale variables and parametrizing the influence of the unresolved small-scale processes on the evolution of the resolved variables in the two-level L96. We consider two stochastic parametrization schemes for the unresolved processes: an empirical parametrization and another one based on response theory. The first one was proposed by Wilks (2005), using multivariate regression and a simple first-order autoregressive model. This parametrization has been widely applied to the two-level L96 and showed very good performance in reproducing the large-scale variables of the full dynamic model Arnold *et al.* (2013); Christensen *et al.* (2015); Harlim (2017). The second parametrization scheme was recently introduced by Vissio and Lucarini (2018), using the methodology proposed by Wouters and Lucarini (2012, 2013, 2016). This is a scale-adaptive parametrization, such that one can derive a universal expression for the parametrization that explicitly represents a scale parameter to do with the coupling strength between the resolved and unresolved processes. The two parametrizations show comparable skills in reproducing the probability density, the spatial correlation, and the temporal autocorrelation of the large-scale variables of the full model Vissio and Lucarini (2018), with the Wilks approach providing marginally better results at the expense of the fact of being ad-hoc and not adaptive. In this paper, we evaluate how well the two parametrization schemes perform in terms of reproducing the extreme value statistics of the full model. The Wilks parametrization is constructed in such a way that typical events are represented as well as possible but there is no a-priori control on how well extreme events are represented. The Wouters-Lucarini parametrization allows, in principle to control the error in the expectation value of any observable, including those relevant for defining the properties of the tails of the distributions. Conversely, the latter approach suffers from being a perturbative one, with the risk of lack of convergence. In an earlier study, Franzke (2012) showed that a reduced order model constructed by systematic stochastic mode reduction strategy has similar extreme value statistics as the full dynamical model for a wide range of time-scale separations. As a contribution to the main goal of this paper, like Franzke (2012); Gálfi *et al.* (2017); Bódai (2017), we also examine the convergence/approximation of the shape parameter estimates of the said different models to the theoretical value of the two-level L96. Note that the shape parameter of a stochastic model is most probably not the same as that of a deterministic model.

In the following section, we will introduce the two-level L96 and the two parametrization schemes. In Sec. 3, we present the two approaches from EVT and the mathematical expression for the shape parameter in dynamical systems. In Sec. 4, we apply the two approaches of EVT

to the parametrized and full models and compare the estimates of the GEV and GP parameters obtained for the different models. In Sec. 5, we provide an empirical comparison of the extremes in the different models. In Sec. 6, we compare the mean return periods of the extremes of the same magnitude in the different models. In Sec. 7, we repeat the analysis for a different coupling strength in the two-level L96. We close the paper with a summary and conclusions in Sec. 8.

2 The two-level Lorenz'96 Model

The two-level L96 was introduced by Lorenz (1995), the governing equations of which describe the dynamics of a lattice with periodic boundary conditions and represent, in a very conceptual way, the main processes occurring in the atmosphere: advection, forcing, and dissipation. In order to apply the Wouters-Lucarini (W-L) parametrization to the two-level L96, Vissio and Lucarini (2018) made two changes to the original model: 1) introduced a forcing term in the equations for the small-scale variables, and 2) restricted the periodic boundary conditions of the small-scale variables within the corresponding large-scale sectors. The first change was aimed at fulfilling a basic requirement for the W-L parametrization: the presence of chaos in the uncoupled dynamics, so that the autocorrelation of the variables decays fast. This requires, in physical terms, an external forcing providing energy to the small-scale variables. The second change was implemented in order for the small-scale variables to represent subgrid-scale phenomena of the sectors they belong to. Additionally, the latter made the implementation of the W-L parametrization easier. The modified governing equations of the two-level L96 are given as:

$$\frac{dX_k}{dt} = -X_{k-1}(X_{k-2} - X_{k+1}) - X_k + F_x - \frac{hc}{b} \sum_{j=1}^J Y_{j,k}, \quad (1)$$

$$\frac{dY_{j,k}}{dt} = -cY_{j+1,k}(Y_{j+2,k} - Y_{j-1,k}) - cY_{j,k} + \frac{c}{b}F_y + \frac{hc}{b}X_k, \quad (2)$$

where the variables X_k and variables $Y_{j,k}$ are defined for $k = 1, \dots, K$ and $j = 1, \dots, J$. We consider the variables X_k to be large-scale variables, while the variables $Y_{j,k}$ to be small-scale variables. The boundary conditions are defined as:

$$\begin{aligned} X_{-1} &= X_{K-1}, \\ X_0 &= X_K, \\ X_{K+1} &= X_1, \\ Y_{J+1,k} &= Y_{1,k}, \\ Y_{J+2,k} &= Y_{2,k}, \\ Y_{0,k} &= Y_{J,k}. \end{aligned}$$

The parameters F_x and F_y represent forcing terms in the equations for the variables X_k and $Y_{j,k}$, respectively. The parameter h is a coupling coefficient, and the parameters c and b can be thought of as time-scale ratio and spatial-scale ratio, respectively. These three parameters determine the coupling strength of the system: $\epsilon = hc/b$. As we keep b and c equal, we have $\epsilon = h$. The variables X_k can represent some atmospheric quantity in K sectors of a latitude cycle, while the variables $Y_{j,k}$ can represent some other atmospheric quantity in smaller JK sectors. There are J smaller sectors in each larger sector. In our computations we set the parameter values of the model as follows: $K = 10$, $J = 10$, $F_x = 10.0$, $F_y = 6.0$, $h = 1.0$, $b = 10.0$, and $c = 10.0$.

In many practical applications, small-scale processes are too expensive to resolve and their impacts on the evolution of large-scale processes are parametrized by deterministic terms, or stochastic terms, or both of them. In the two-level L96 we can parametrize the effects of the evolution of the $Y_{j,k}$ variables on the evolution of the X_k variables, then the evolution equations of X_k variables are given as

$$\frac{dX_k}{dt} = -X_{k-1}(X_{k-2} - X_{k+1}) - X_k + F_x + U, \quad (3)$$

where U denotes the parametrization of the effects of the unresolved processes, which has to represent the model error when only the large-scale variables X_k are resolved in place of the full dynamics. We call Eq. (3) the parametrized model, contrasting with the full model given by Eqs. (1) and (2). We now introduce two parametrization schemes.

2.1 Wilks parametrization: multivariate regression and autoregressive process

The first parametrization scheme was proposed by Wilks (2005), which used a polynomial equation and a correlated noise term to represent the unresolved processes:

$$U(t) = P_k(X_k(t)) + e_k(t), \quad (4)$$

where the polynomial equation is a function of the X_k variables:

$$P_k(X_k(t)) = a_0 + a_1 X_k(t) + a_2 X_k^2(t) + a_3 X_k^3(t) + a_4 X_k^4(t), \quad (5)$$

and the noise term is a simple first-order autoregressive model:

$$e_{k,i} = \phi e_{k,i-1} + \sigma_e (1 - \phi^2)^{1/2} z_{k,i}, \quad (6)$$

where $e_{k,i} = e_k(t_i)$, $t_i = dt \times i$, $i = 1, \dots, I$, $dt = 0.005$ [MTU], where MTU is the abbreviation of ‘model time unit’, and $z_{k,i}$ are unit variance Gaussian random variables, independent wrt. both k and i . The autoregressive parameter ϕ is equal to the autocorrelation of the time series $e_k(t_i)$ with a time lag of dt , and σ_e denotes the standard deviation of the time series. We first estimate the polynomial coefficients by regressing $P_k(X_k(t_i))$ against

$$-\frac{hc}{b} \sum_{j=1}^J Y_{j,k}(t_i).$$

The estimated coefficients are $a_0 = -1.81$, $a_1 = -0.1467$, $a_2 = 0.001357$, $a_3 = -0.001446$, and $a_4 = 0.0001313$. Next we fit the residual of the polynomial fitting

$$-\frac{hc}{b} \sum_{j=1}^J Y_{j,k}(t_i) - P_k(X_k(t_i); a_0, a_1, a_2, a_3, a_4)$$

with the first-order autoregressive model (Eq. 6). In our computation, we get $\phi = 0.9997$ and $\sigma_e = 0.8965$; we refer to Neumaier and Schneider (2001) for the estimation of parameters of autoregressive models. The Wilks parametrization is an empirical parametrization which is constructed based on the fact that in the two-level L96 the unresolved tendency is strongly and nonlinearly dependent on the value of the resolved variable (Wilks, 2005). The implementation of the two modifications to the original model do not change this characteristic, therefore, the Wilks parametrization is still valid. A weakness of the empirical parametrizations is that their parameters need to be recalculated if the configuration of the full model is changed.

2.2 Wouters–Lucarini parametrization: averaging, correlations and memory

The application of the second parametrization scheme to the two-level L96 was first demonstrated by Vissio and Lucarini (2018). The W-L parametrization is based on Ruelle’s response theory (Ruelle, 1997, 2009), and it was proposed by Wouters and Lucarini (2012). Later Wouters and Lucarini (2013) showed that the parametrization scheme can also be obtained through the Mori–Zwanzig approach Mori et al. (1974); Zwanzig (1960, 1961). Note that this point of view, as well as the more sophisticated results by Chekroun, Liu, and Wang (2015), provide the mathematical foundations of the theory of stochastic parametrizations. In the W-L parametrization, the coupling of the variables to be parametrized is considered as a small perturbation to the variables of interest. The coupling is decomposed into three terms: an averaging term, a correlation term, and a memory term. Therefore, the formula of the W-L parametrization is given as:

$$U = D + S_k + M_k. \quad (7)$$

In the above the averaging term D is a constant, accounting for the ‘‘averaged influence’’ of the $Y_{j,k}$ variables on the long-term statistics of the X_k variables, and it is calculated by

$$D = -\frac{1}{b} \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \sum_{j=1}^J \tilde{Y}_{j,k}(\tau) d\tau, \quad (8)$$

where $k = 1, \dots, K$. Due to symmetry, this term has the same value for all X_k variables. The time series of $\tilde{Y}_{j,k}$ are obtained by integrating the equations given as:

$$\frac{d\tilde{Y}_{j,k}}{d\tau} = -\tilde{Y}_{j+1,k}(\tilde{Y}_{j+2,k} - \tilde{Y}_{j-1,k}) - \tilde{Y}_{j,k} + F_y, \quad (9)$$

where the $\tilde{Y}_{j,k}$ denote the rescaled small-scale variables:

$$\tilde{Y}_{j,k} = bY_{j,k}, \quad (10)$$

and time is rescaled as

$$\tau = ct. \quad (11)$$

The correlation term S_k is a stochastic term which accounts for the fluctuations of the influence of the $Y_{j,k}$ variables on the long-term statistics of the X_k variables. It is constructed as an additive noise which reproduces the temporal correlation of the fluctuations. The autocovariance of the fluctuations is given as:

$$R_k(\tau) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T p_k(\tau_1) p_k(\tau_1 + \tau) d\tau_1, \quad (12)$$

where

$$p_k(\tau) = - \sum_{j=1}^J \frac{\tilde{Y}_{j,k}(\tau)}{b} - D. \quad (13)$$

Following Vissio and Lucarini (2018), we generate the stochastic term S_k using a simple autoregressive model. The M_k is a non-Markovian term, accounting for the memory effects, which is important for the parametrization of the small-scale processes of the two-level L96; see the comparison of the first- and second-order parametrizations in Vissio and Lucarini (2018), where the first one contains only the averaging term, while the second one contains also the correlation term and the memory term. M_k describes the influence of the past values on the present values of the X_k variables through the coupling of the $Y_{j,k}$ variables, and an explicit expression for M_k was provided by Vissio and Lucarini (2018):

$$M_k = -\frac{1}{b} \int_0^\infty X_k(t - \tau_1) H(\tau_1) d\tau_1, \quad (14)$$

where

$$H(\tau_1) = \frac{1}{\Omega} \lim_{\Omega \rightarrow \infty} \int_0^\Omega \sum_{j=1}^J \frac{\partial \tilde{Y}_{j,k}(\tau_1 + \omega)}{\partial \tilde{Y}_{j,k}(\omega)} d\omega. \quad (15)$$

The W-L parametrization is a scale-adaptive parametrization; in the two-level L96, it can be analytically adapted to the changes in the values of h , b , and c by simply rescaling the three terms D , S_k and M_k . However, at the expense of being flexible, it may be somewhat less accurate than the empirical parametrization. A limitation of the W-L parametrization is that it is only valid for weakly coupled systems.

2.3 Local and Global Observables

In physics an observable is a variable that can be measured. We consider a local observable and two global observables of the full and parametrized models; we define the local observable as $A_x = X_k$, for $k = 1, \dots, K$ (the X_k variables have the same statistics due to the symmetry of the model), and the global observables as $A_E = \sum_{k=1}^K X_k^2$ and $A_p = \sum_{k=1}^K X_k$, representing the total energy and the total momentum of the X_k variables, respectively. Fig. 1 compares the probability density functions (PDF) of the local observable between the full and parametrized models, and Fig. 2 compares the PDFs of the global observables between the full and parametrized models. The sample data of the local and global observables are produced by integrating the models for 3.2×10^5 time steps with adaptive stepsizes. We record the three observables at each time step and since the 10 different X_k variables are statistically equivalent, we get 3.2×10^6 samples for the local observable. The two parametrized models can capture the statistics of the considered observables of the full model relatively well, but they produce more precise statistics for the local observable than the global observables. Moreover, compared to the W-L parametrized model, the Wilks parametrized model better reproduces the statistics of the local observables of the full model. However, it must be noted that what we look at here is the bulk of the statistics, and these figures do not allow us to compare the extreme value statistics.




Figure 1: Probability density functions of the observable A_x from the full and parametrized models.




Figure 2: Probability density functions of the observables A_E (left) and A_p (right) from the full and parametrized models.

3 Extreme Value Analysis

3.1 Two Approaches

Two straightforward approaches exist for extreme value analysis: the BM approach and the POT approach; we refer the reader to Embrechts, Klüppelberg, and Mikosch (1997), Coles (2001), and Lucarini *et al.* (2016) for more details. According to the BM approach the data is divided into sufficiently large blocks of equal length and the maximal value of each block is retained. Then the GEV distribution is fitted to these maxima. The cumulative distribution function (CDF) of the GEV distribution is given as:

$$P(X \leq x) = F(x; \mu, \sigma, \xi) = \exp \left\{ 1 - \left[1 + \xi \left(\frac{x - \mu}{\sigma} \right) \right]^{-1/\xi} \right\}, \quad (16)$$

which holds for $1 + \xi(x - \mu)/\sigma > 0$, where $-\infty < \mu < \infty$ is referred to as location parameter, $\sigma > 0$ is called scale parameter and $-\infty < \xi < \infty$ denotes the so-called shape parameter. When $\xi = 0$, we take the limit of (16) as $\xi \rightarrow 0$, which is

$$F(x; \mu, \sigma) = \exp \left\{ -\exp \left[- \left(\frac{x - \mu}{\sigma} \right) \right] \right\}. \quad (17)$$

The location and scale parameters are scaling constants used to normalize the random variable X of the maximal values of blocks. The shape parameter characterises the tail behavior: when $\xi = 0$, the tail decays exponentially; when $\xi > 0$, a heavy tail occurs, which decays following a power law; and when $\xi < 0$, the tail is bounded, i.e., there is an upper limit of the domain of the distribution. The POT approach selects data whose magnitude is above a high threshold. The threshold exceedances are fitted by the GP distribution, which is defined by the CDF:

$$P(\hat{X} \leq \hat{x}) = F(\hat{x}; \hat{\sigma}, \hat{\xi}) = 1 - \left(1 + \hat{\xi} \frac{\hat{x}}{\hat{\sigma}}\right)^{-1/\hat{\xi}}, \quad (18)$$

where $\hat{X} > 0$ denotes the exceedances above a threshold u , e.g. $\hat{X} = X_k - u$. The scale parameter $\hat{\sigma} > 1$, and $1 + \hat{\xi}\hat{x}/\hat{\sigma} > 0$. The shape parameter $\hat{\xi}$ again characterises the tail behaviour like it does in the GEV distributions. As with the GEV distribution, for $\hat{\xi} = 0$, the CDF is given as:

$$F(\hat{x}; \hat{\sigma}) = 1 - \exp\left(-\frac{\hat{x}}{\hat{\sigma}}\right). \quad (19)$$

It should be noted that the shape parameter is the same for corresponding limit GEV and GP distributions, i.e., $\xi = \hat{\xi}$ Lucarini et al. (2016). A simple functional relation connects the two distributions; the natural logarithm of Eq. (16) plus one equals Eq. (18). Under general conditions, while the two approaches define extremes differently, they are fundamentally equivalent.

3.2 The Theoretical Value of Shape Parameter

For general classes of smooth physical observables, Holland *et al.* (2012) and Lucarini *et al.* (2014) provided an explicit expression for the shape parameter:

$$\xi_{theo} = -\frac{1}{\delta}, \quad (20)$$

where

$$\delta = d_s + \frac{(d_u + d_n)}{2}, \quad (21)$$

for continuous-time flows. In the above, d_u is equal to the number of positive Lyapunov exponents, d_n is equal to the number of zero exponents, which is at least one for Axiom A systems (Lucarini et al., 2014), and d_s equals the number of stable directions and is given by

$$d_s = d_{KY} - d_u - d_n. \quad (22)$$

In the above, the Kaplan-Yorke dimension is defined as Kaplan and Yorke (1979)

$$d_{KY} = n + \frac{\sum_{i=1}^n \lambda_i}{|\lambda_{n+1}|}, \quad (23)$$

where λ_i denote the Lyapunov exponents of the system, arranged in a descending order, and n is the number when $\sum_{i=1}^n \lambda_i$ is larger than zero while $\sum_{i=1}^{n+1} \lambda_i$ is smaller than zero.

Eqs. (20-23) give an estimator for the shape parameter via the estimators of the dimensions, or via the estimators of the Lyapunov exponents, instead of estimating it by fitting a GEV or a GP distribution. Clearly, the value of ξ_{theo} is always negative; this can be seen from the formulae, where we always have $\delta > 0$. A negative shape parameter indicates that the distribution of extremes has an upper bound. This is not surprising as smooth physical observables are considered and the dynamical attractor is compact Lucarini et al. (2014). Finite-size positive estimates are of course possible. We should think that whenever positive shape parameters are and have been encountered, a deterministic dynamical system is still an appropriate mathematical representation of the physical process. When we consider a high-dimensional chaotic system, we have a large Kaplan-Yorker dimension of the attractor and so the value of ξ_{theo} is close to zero. It means that the occurrence of very large extreme events becomes more likely.




Figure 3: Comparison of the estimates of the GEV shape parameter for the observable A_x (left), A_E (middle), and A_p (right) over a range of block sizes. The horizontal line shows the theoretical value of the shape parameter for the full model.

3.3 Return Period

The mean or expected return period is the reciprocal of the probability of an extreme value of a given magnitude. The magnitude of the extreme value in this context is commonly referred to as the return level. The fitted GEV and GP distributions provide estimates of mean return periods for different return levels via inverting those distributions. As an extrapolation, one can choose higher return levels than ever observed, however, the corresponding mean return time estimates are in general biased. We can calculate the empirical return periods for the data up to reasonable return levels, which can be compared – in order to facilitate a goodness-of-fit check – to the said estimates by fitting a GEV distribution and inverting it. Using the block maxima data, the empirical mean return period for return level m is calculated as:

$$\tau_{BM}(m) = \left(\frac{n_p(m)}{n_b} \right)^{-1}, \quad (24)$$

where n_p is the number of block maxima which have a value greater than or equal to m , and n_b is the number of blocks, or the size of data. We construct mean return time vs return level plots to compare the full and parametrized models yet another way. We create these diagrams using the package "extRemes 2.0" Gilleland and Katz (2016) of the software environment R.

4 Comparison of GEV and GP Parameters

We apply the BM and POT methods to the local and global observables of the full and parametrized models. We consider a range of block sizes B , exponentially increasing from the smallest block size considered $B_0 = 1000$ MTU, and a range of exceedance ratios E , exponentially decreasing from the highest ratio considered $E_0 = 0.1\%$, for the BM method and POT method, respectively. In order to compare fairly the GP and GEV parameters estimated from the samples of the same size, we match choices of B with choices of $E = 1/B$. The number of the block maxima given by the block size B_0 , equaling the number of the threshold exceedances given by the exceedance ratio E_0 , is 4×10^6 for the local observable (4×10^5 for each X_k variable) and 4×10^5 for the global observables. We take advantage of the package "extRemes 2.0" Gilleland and Katz (2016) of the software environment R to estimate the GEV and GP parameters; we use the method of L-moments Hosking (1990) for estimation. For the local observable, we plot the averaged estimates of the parameters over ten X_k variables, and present the estimator variance as the sample standard deviation for the ten estimates given by error bars. For the global observables on the other hand, we show the 95% confidence intervals of each estimate, where the confidence interval is calculated by a parametric bootstrap method Gilleland and Katz (2016).

4.1 The GEV Parameters

Fig. 3 compares the estimates of the GEV shape parameter for the three observables of the full and parametrized models over a range of block sizes. The horizontal line represents the theoretical




Figure 4: Comparison of the estimates of the GEV scale parameter for the observable A_x (left), A_E (middle), and A_p (right) over a range of block sizes.




Figure 5: Comparison of the estimates of the location parameter for the observable A_x (left), A_E (middle), and A_p (right) over a range of block sizes.




Figure 6: Comparison of the estimates of the GP shape parameter for the observable A_x (left), A_E (middle), and A_p (right) over a range of exceedance ratios. The horizontal line shows the theoretical value of the shape parameter for the full model.




Figure 7: Comparison of the estimates of the GP scale parameter for the observable A_x (left), A_E (middle), and A_p (right) over a range of exceedance ratios.

value of the shape parameter for the full model, which is not applicable to the parametrized models, because they obey a stochastic and not deterministic dynamics. The shape parameter determines the tail behaviour of the distribution: a larger value indicates a slower decay of the tail; and, on the contrary, a smaller value indicates a faster decay of the tail.

We first focus on the behaviour of estimates of the shape parameter for the full model. We observe nonmonotonic changes of the estimates with an increase of block size. The estimates for the observable A_x decrease with the block size increasing from B_0 to $B_0 \times 2^3$, then it turns to increase and is seen to cross the theoretical value at $\log_2(B/B_0) = 7$, and after that the theoretical value is always within the increasingly large confidence intervals of the estimates. The estimates for the observable A_E cross the theoretical value at $\log_2(B/B_0) = 4$ and then seems to come back at $\log_2(B/B_0) = 11$, however, the confidence interval is very big by then. The estimate for the observable A_p increases at the first three data points, and then the estimates are always within the confidence intervals of any latter estimates. The estimates for the observable A_p seem to reach a value smaller than the theoretical value, however, it should be noted that the last three data points show an upward trend, toward the theoretical value. As the block size increases, we observe to some extent an approximation of the estimates of the shape parameter to the theoretical value instead of a steady convergence. The estimates for the three observables show different erratic behaviours with the increase of the block size, and so none of these behaviours is asymptotic.

We now focus on comparing the estimates of the shape parameter given by the parametrized models to that given by the full model. When we look at the observable A_x and A_E , the estimates given by the parametrized models are apparently different from that given by the full model over most of the block sizes examined. An interesting result is that when we look at the observable A_p , the W-L parametrized model gives very similar estimates compared to the full model. We note, however, that this good correspondence could be just a coincidence; the results may change if we use another system or a different parameter setting. Moreover, this good correspondence should not last as the stochastic model surely has a different asymptotic shape parameter from that of the deterministic model.

Fig. 4 compares the estimates of the GEV scale parameter for the three observables of the full and parametrized models over a range of block sizes. The estimates for the observable A_x and A_p given by all the three models monotonically decrease as the block size increases, while the estimate for the observable A_E nonmonotonically changes. The parametrized models generally give larger estimates than the full model for all the three observables over all the block sizes considered. We observe that the W-L parametrized model gives more accurate estimates for the local observable over all considered block sizes than the Wilks parametrized model. Similarly to the estimates of the shape parameter for the observable A_p , this result could also be just a coincidence.

Fig. 5 compares the estimates of the location parameter for the three observables from the full and parametrized models over a range of block sizes. The location parameter, only appearing in GEV fittings, determines where the center of the distribution is located, and a larger value means that the distribution is shifted to the right, so that we have extremes of higher magnitudes. We observe that both of the parametrized models give larger estimates for all the three observables than the full model. Additionally, the W-L parametrized model gives more accurate estimates for the observable A_x and A_E than the Wilks parametrized model.




Figure 8: Comparison of the histograms of the block maxima of the observable A_x (left panels), A_E (middle panels), and A_p (right panels) between the full and W-L parametrized models (upper panels), and between the full and Wilks parametrized models (lower panels). The block size is B_0 .

4.2 The GP Parameters

Fig. 6 compares the estimates of the GP shape parameter for the three observables of the full and parametrized models over a range of exceedance ratios. Fig. 7 compares the estimates of the GP scale parameter for the three observables of the full and parametrized models over a range of exceedance ratios. The behaviours of the estimates of the GP parameters are similar to those of the GEV parameters. We shall not comment on these further. The BM and POT approaches give consistent results of extreme value statistics as they are fundamentally equivalent.

5 Empirical Comparison of Extremes

5.1 Histograms of Block Maxima

Fig. 8 compares the histograms of the block maxima of the three observables between the parametrized and full models. The total area of each histogram is normalized to one, and so the vertical axes of the histograms show relative probability density. The histograms give a rough estimate of the PDF of the block maxima. Since we have a large number of data points and small bin sizes, the graphs look very smooth. The patterns of the histograms are unimodal and approximately symmetric; they are slightly right skewed like indeed a GEV distribution. The disagreements between the histograms of the parametrized models and the full model have two main features: 1) the block maxima from the parametrized models are more widely distributed; and 2) the histograms of the parametrized models are right-shifted. These empirical results have been already indicated by the estimates of the scale and location parameter in Sec. 4. Fig. 9 shows the same figures as Fig. 8 but with the block maxima selected by a larger block size. As the block size increases, the block maxima have larger magnitudes and the overlapping area between the histograms of the parametrized and full models becomes smaller. If we look at the observable A_E , we can see that the probability density of the block maxima given by the Wilks parametrized model is totally different from that given by the full model. In summary, the histograms of the W-L parametrized model appear to be in somewhat better agreement with the histograms of the full model than that of the Wilks parametrized model. Moreover, the local observable of the full and parametrized models show better agreement than the global observables.




Figure 9: Comparison of the histograms of the block maxima of the observable A_x (left panels), A_E (middle panels), and A_p (right panels) between the full and W-L parametrized models (upper panels), and between the full and Wilks parametrized models (lower panels). The block size is $B_0 \times 2^7$.

5.2 Number of Threshold Exceedances

When applying the POT method, a main issue is to compare the number of the threshold exceedances produced by the full and parametrized models for a given threshold. Fig. 10 shows the fractional difference of the number of the threshold exceedances given by the parametrized models compared to that given by the full model. We calculate the fractional difference using the formula given as:

$$\frac{n_{par}(u) - n_{full}(u)}{n_{full}(u)}, \quad (25)$$

where $n_{full}(u)$ denotes the number of exceedances above a threshold u from the full model, and $n_{par}(u)$ denotes that from the parametrized models. We consider two sets of thresholds for the three observables of the three models. The first set of thresholds is chosen as the 99.9th percentile of each observable of the full model, and the second set of thresholds is more stringent, which is the 99.99th percentile of each observable of the full model. As shown in Fig. 10, both of the parametrized models give more threshold exceedances than the full model, with the exception of the observable A_E of the W-L parametrized model in the upper panel. Moreover, the global observables demonstrate larger differences than the local observable with the same exception. When considering the observable A_x and A_E , the W-L parametrized model has a substantially better performance than the Wilks' one. Note that if we look at the observable A_E of the Wilks parametrized model in the lower panel, we can see a huge discrepancy; the number of threshold exceedances given by the Wilks parametrized model is more than 11 times greater than that given by the full model. Instead, the performance of the two parametrizations is comparable when looking at the observable A_p . We also observe that the fractional difference becomes larger when a higher threshold is given.

5.3 Quantile-Quantile Plots

We further compare the empirical quantiles of the three observables from the parametrized models to those from the full model. We compare 1001 quantiles, starting with the smallest element of the sample, up to the largest element, and 999 equidistant quantiles in between these two




Figure 10: The fractional difference of the number of threshold exceedances from the parametrized models compared to that from the full model. In the upper (lower) panel, the threshold is the 99.9th (99.99th) percentile of the corresponding observable of the full model.




Figure 11: The empirical quantiles of the observable A_x (left), A_E (middle), and A_p (right) from the parametrized models against that from the full model.

elements. Fig. 11 shows the empirical quantiles of the three observables from the parametrized models against those from the full model. In the left panel, except for several very high-level and low-level quantiles, the other moderate quantiles from the parametrized models agree well with those from the full model, indicating that the parametrized model can reproduce well the bulk statistics of the observable A_x of the full model. In the middle panel, the points are slightly farther away from the diagonal, indicating that the quantiles of the observable A_E from the parametrized models are somewhat in disagreement with those from the full model. In the right panel, we observe that the quantiles of the observable A_p from the Wilks parametrized model agree well with those from the full model, whereas the low quantiles from the W-L parametrized model show an obvious disagreement. In summary, the parametrized models give better quantiles of the local observable than the global observables.

The data between the last two quantiles are the largest one thousandth of the sample data. These data are the extreme values we want to investigate. Fig. 12 shows a zooming onto the extreme quantiles that are higher than the 99.9th percentile. We consider 101 extreme quantiles, including the 99.9th percentile, the largest element of the sample, and 99 equidistant quantiles between them. As shown in the left panel, the extreme quantiles of the observable A_x from the W-L parametrized model are slightly larger than those from the full model, and, in comparison,




Figure 12: The empirical extreme quantiles of the observable A_x (left), A_E (middle), and A_p (right) from the parametrized models against that from the full model.

the extreme quantiles from the Wilks parametrized model have a larger disagreement. In the middle panel, the extreme quantiles of the observable A_E from the Wilks parametrized model totally miss that from the full model and the approximation of the extreme quantiles from the W-L parametrized model to those from the full model breaks down after the 99.95th percentile. In the right panel, we observe that the extreme quantiles of the observable A_p from both of the parametrized models are far away from those from the full model.

6 Return Periods of Extremes

Fig. 13 presents the mean return periods against the return levels of the extremes for the three observables of the three models. The extremes are selected by the BM approach using a block size of $B_0 \times 2^{10}$. It can be seen that the GEV mean return time estimates using the chosen block size fit the empirical data quite well, except for very high return levels of the observable A_x , at which the empirical estimates are out of the confidence interval of the GEV estimates.

Comparing the parametrized models to the full model, we find that for all the three observables the parametrized models always produce higher-level extremes in the same return periods. In other terms the return period of the extremes with the same magnitude is shorter in the parametrized models than in the full model. Moreover, in comparison with the Wilks parametrized model, the W-L parametrized model gives better estimates of return levels, which are closer to the estimates given by the full model. One exception is observed in momentum observable; the Wilks parametrized model gives slightly better estimates of return levels than the W-L parametrized model.

7 Sensitivity to coupling strength

In previous numerical experiments, we applied the Wilks and W-L parametrizations to the two-level L96 with the standard value of the coupling strength $h = 1.0$. We now test and compare the performances of the two parametrizations as we consider both a factor of 2 weaker and a stronger coupling strengths, i.e., $h = 0.5$ and 2 , respectively.

We need to recompute the parameter values for the Wilks parametrization through fitting the residual time series as we vary the coupling strength, but this is not the case for the W-L parametrization, the three terms of which can be simply rescaled as follows (see Eq. (13) of Vissio and Lucarini (2018)): $\tilde{D} = (\tilde{h}/h)D$, $\tilde{S} = (\tilde{h}/h)S$ and $\tilde{M} = (\tilde{h}/h)^2 M$, where the symbols with tilde represent the case when we have a new value of h , and the symbols without tilde represent the case when we have the standard value of the coupling strength. The parameter values of the Wilks parametrization for the weaker and stronger coupling strength are presented in Table 1. Note that the Wilks parametrization lacks stability in the case of $h = 2.0$ if we include the stochastic term in the parametrization, therefore, we use a deterministic parametrization, which does not contain the first-order autoregressive model.

Fig. 14 shows the PDF of the three observables from the full and parametrized models in the case of $h = 0.5$. We observe a better agreement for all the three observables between the full and parametrized models in comparison with the case of $h = 1.0$. Because of the way that the




Figure 13: Return levels against return periods for the block maxima of the observable A_x (upper panels), A_E (middle panels), and A_p (lower panels) from the full and parametrized models. The black dots signify the empirical return levels, the solid black lines show the predicted return levels computed based on the fitted GEV distributions, and the dashed grey lines present the 95% confidence intervals of the estimates.




Figure 14: Probability density functions of the observable A_x (left), A_E (middle), and A_p (right) from the full and parametrized models in a weak coupling case ($h = 0.5$).

Table 1: Parameter values of the Wilks parametrization for different values of the coupling strength h .

h	a_0	a_1	a_2	a_3	a_4	ϕ	σ_e
0.5	-0.95	-0.0349	-0.0008	-0.0002	0.00002	0.967	0.3548
2.0	-3.54	-0.3292	-0.0548	-0.0593	0.0132	0.0	0.0

W-L parametrization is constructed, the error in the W-L parametrized model with respect to the full model is a function of the coupling strength, which is $\mathcal{O}(\epsilon^3)$, as mentioned before, and this applies to all observables in the system Lucarini et al. (2014); Vissio and Lucarini (2018). Therefore, it is unsurprising that the W-L parametrized model demonstrates a better performance in the case of a weaker coupling strength. The Wilks parametrization is constructed based on the correlation between the X_k variables and the coupling term, which is the sum of the $Y_{j,k}$ variables. As the coupling strength becomes weaker, this correlation also becomes smaller, however, in the meanwhile, the difference between the uncoupled dynamics and the coupled dynamics reduces, therefore, although the quality of the parametrization scheme is not automatically constrained to improve, we still observe smaller difference between the parametrized model and the full model.

Fig. 15 shows the PDF of the observable A_x from the full and parametrized models in the case of $h = 2.0$. In this case, the statistics of the observable A_x in the full model are greatly changed; the PDF changes from unimodal to multimodal. The Wilks parametrization is able to capture the multimodality of the PDF, since the correlation between the X_k and the $Y_{j,k}$ variables still exist. However, we believe that this multimodality prevents us to see a cubic scaling for the difference between the full and W-L parametrized models. The qualitative difference shown in the PDFs suggests that we are outside of a radius of convergence with respect to the coupling strength $h = 2.0$, which means that the scale adaptivity is only valid for a range of the coupling strength. This is unsurprising, because the W-L parametrization can only be applied to weakly coupled systems. The PDFs for the two global observables, A_E and A_p , have not been shown, because both parametrizations perform very poorly.

We now look at the extreme value statistics given by the two parametrized models in the case of $h = 0.5$. Fig. 16 compares the estimates of the GP shape parameter for the three observables between the full and parametrized models. In the case of a weaker coupling, the parametrized models display extreme value statistics that matches more closely that of the full model in comparison with the case of the original coupling strength. Furthermore, we compare the number of threshold exceedances for the parametrized models to those for the full model. The thresholds are chosen as the 99.9th and 99.99th percentile of the data produced by the full model. Fig. 17 shows the fractional difference calculated by Eq. (25) in a weak coupling case. The W-L parametrized model gives a closer number of the extremes of the observable A_x and A_E to the full model than the Wilks parametrized model, and the Wilks parametrized model gives a closer number of the extremes of the observable A_p . These are features unchanged with respect to the standard parameter setting. Otherwise, both the parametrized models give less extremes of the observable A_p than the full model.




Figure 15: Probability density functions of the observable A_x from the full and parametrized models in a strong coupling case ($h = 2$).




Figure 16: Comparison of the estimates of the GP shape parameter for the observable A_x (left), A_E (middle), and A_p (right) over a range of exceedance ratios in a weak coupling case ($h = 0.5$).



Figure 17: The fractional difference of the number of threshold exceedances from the parametrized models compared to that from the full model. In the upper (lower) panel, the threshold is the 99.9th (99.99th) percentile of the corresponding observable of the full model. Coupling strength $h = 0.5$.

8 Summary and Conclusions

This paper addresses the problem of how well stochastic parametrizations perform, in a simple yet informative case; we conduct numerical experiments in a conceptual atmospheric model, the two-level L96 Lorenz (1995). Instead of investigating the bulk of the statistics, as what has been usually done, here we focus on examining extreme value statistics. The problem has been explored using EVT, which provides us with a mature statistical framework to define and analyze extremes. Under this framework, probability distributions obey a limit law and so some universality, as opposed to the bulk statistical properties. We considered two kinds of parametrization schemes for the two-level L96: the widely applied Wilks parametrization scheme Wilks (2005) and a recently proposed scale-adaptive parameterization scheme Vissio and Lucarini (2018). The latter one is constructed by the methodology proposed by Wouters and Lucarini (2012, 2013, 2016) and hence denoted as W-L parametrization. The parametrized models constructed by these two schemes can reasonably reproduce the bulk, but not the extreme value statistics of the large-scale variables of the full model. The differences are mainly found in the following aspects:

- The extreme values from the parametrized models have different shape parameters, scale parameters and location parameters than those from the full model;
- The extreme values selected by the BM approach have a larger variance in the parametrized models than in the full model;
- The parametrized models give different numbers of extreme events than the full model when we apply the same threshold to them, and in most cases, the parametrized models produce more extreme events.

From a practical point of view, these differences will affect the mean return times of extreme events; the extreme events from the parametrized models have typically shorter return times than those from the full model. Our results suggest that stochastic parametrizations should be accurately tested against their performance on extremes, because a good performance on typical events cannot ensure a good performance on untypical events.

The parametrization schemes we use are constructed in two fundamentally different ways, and so they have different advantages and disadvantages. As discussed in Sec. 2, the Wilks parametrization is an empirical one, which is constructed based on the observed fact that in the two-level L96 the small-scale tendency strongly and nonlinearly depends on the value of the large-scale variables Wilks (2005). The W-L parametrization is constructed through a statistical mechanical approach, where the small-to-large-scales coupling is treated as perturbations of strength ϵ added to the uncoupled large-scale dynamics. It is based on a series expansion of the invariant measure of the uncoupled system with respect to the small coupling parameter ϵ . Including first- and second order terms in the W-L parametrization implies that the difference between the expectation value of all observables of the large-scale variables in the full model and in the parametrized model scales as $\mathcal{O}(\epsilon^3)$. In fact, this scaling applies to tail probabilities of our interest too, because they are observables (Lucarini et al., 2014, 2016).

As shown by Vissio and Lucarini (2018), the Wilks parametrization is more accurate than the W-L parametrization in representing the bulk of the statistics of the full model. This is not surprising because the Wilks parametrization is constructed through an empirical fitting procedure that aims at providing an optimal representation of the typical events of the system, Note that this does not ensure a good representation of rare events. Indeed, our new results show that there are many cases where the W-L parametrized model reproduces the extreme value statistics better than the Wilks parametrized model. Figs. 9, 10, 11 are especially clear in this regard. However, we must emphasize that the performance of any parametrization scheme should depend on the observable of interest via a prefactor for the scaling $\mathcal{O}(\epsilon^3)$ of errors. We are yet to develop an understanding why the W-L scheme outperforms Wilks with respect to the local observable A_x of the site variables and the global energy A_E observable but not the global momentum A_p observable.

An important advantage of W-L that sets it apart from Wilks is that if an observable of interest is not reconstructed accurately enough for a purpose, then one can opt – within the perturbative framework – to work out higher-order correction terms. Within a radius of convergence of the series expansion with respect to ϵ , the performance can be improved arbitrarily in principle. The other major advantage of the W-L parametrization is the scale adaptivity; when the coupling strength ϵ is changed, we can adapt the obtained parametrization schemes by a suitable rescaling

procedure. In contrast, we need to recompute the parameter values of the Wilks parametrization every time when the coupling strength is changed. But if the coupling strength is too strong, the W-L parametrization fails; it is outside of the radius of convergence, while the Wilks parametrization might still work (although even that was found to be not the case for $h = 2$, only the deterministic parametrization). This shows advantages and disadvantages of the two strategies.

In this paper, we also examined the convergence (or approximation, with a finite residual error possibly) of finite size estimates of the GEV and GP shape parameters to the theoretical value of the deterministic model, which is computed by the partial dimensions of the attractor of the system. In the two-level L96, the estimates of the shape parameter non-monotonically approach the theoretical value as the block size or threshold increases. The estimates show erratic behaviours rather than an asymptotic behaviour in the observed ranges of the block maxima and thresholds. The approximation of the estimates to the theoretical value takes place very slowly, and they have different rates for different observables. Different rates of convergence of the estimates of the shape parameter for different observables had also been reported in Gálfi, Bódai, and Lucarini (2017).

Statement on computer code availability: All code used for this study and written by us is available on request.

Acknowledgments

GH is funded by China Scholarship Council (CSC). TB and VL acknowledge financial support from the Horizon 2020 projects CRESCENDO (under grant No. 641816) and Blue-Action project (under grant No. 727852). We thank Gabriele Vissio for providing the code for the W-L parametrized model. We thank two anonymous reviewers for their effort with reviewing our manuscript and for providing feedback.

References

- H. M. Arnold, I. M. Moroz, and T. N. Palmer. Stochastic parametrizations and model uncertainty in the Lorenz '96 system. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 371(1991), 2013. ISSN 1364-503X. doi: 10.1098/rsta.2011.0479. URL <http://rsta.royalsocietypublishing.org/content/371/1991/20110479>.
- A. A. Balkema and L. de Haan. Residual life time at great age. *Ann. Probab.*, 2(5):792–804, 10 1974. doi: 10.1214/aop/1176996548. URL <https://doi.org/10.1214/aop/1176996548>.
- Judith Berner, Ulrich Achatz, Lauriane Batte, Lisa Bengtsson, Alvaro de la Cámara, Hannah M. Christensen, Matteo Colangeli, Danielle R. B. Coleman, Daan Crommelin, Stamen I. Dolaptchiev, Christian L. E. Franzke, Petra Friederichs, Peter Imkeller, Heikki Järvinen, Stephan Juricke, Vassili Kitsios, François Lott, Valerio Lucarini, Salil Mahajan, Timothy N. Palmer, Cécile Penland, Mirjana Sakradzija, Jin-Song von Storch, Antje Weisheimer, Michael Weniger, Paul D. Williams, and Jun-Ichi Yano. Stochastic parameterization: Toward a new view of weather and climate models. *Bulletin of the American Meteorological Society*, 98(3):565–588, 2017. doi: 10.1175/BAMS-D-15-00268.1. URL <https://doi.org/10.1175/BAMS-D-15-00268.1>.
- Tamás Bódai. Extreme Value Analysis in Dynamical Systems: Two Case Studies. In Christian L. E. Franzke and T. J. O’Kane, editors, *Nonlinear and Stochastic Climate Dynamics*, chapter 14, pages 392–429. Cambridge University Press, Cambridge, 2017.
- Mickaël D Chekroun, Honghu Liu, and Shouhong Wang. *Approximation of stochastic invariant manifolds: stochastic manifolds for nonlinear SPDEs I*. Springer, 2015.
- H. M. Christensen, I. M. Moroz, and T. N. Palmer. Simulating weather regimes: impact of stochastic and perturbed parameter schemes in a simple atmospheric model. *Climate Dynamics*, 44(7):2195–2214, Apr 2015. ISSN 1432-0894. doi: 10.1007/s00382-014-2239-9. URL <https://doi.org/10.1007/s00382-014-2239-9>.
- S. Coles. *An Introduction to Statistical Modeling of Extreme Values*. Springer, 2001.

- P Collet. Statistics of closest return for some non-uniformly hyperbolic systems. *Ergodic Theory and Dynamical Systems*, 21(2):401–420, 2001.
- Paul Embrechts, Claudia Klüppelberg, and Thomas Mikosch. *Modelling extremal events for insurance and finance*. Springer Verlag, Berlin, 1997.
- R. A. Fisher and L. H. C. Tippett. Limiting forms of the frequency distribution of the largest or smallest member of a sample. *Mathematical Proceedings of the Cambridge Philosophical Society*, 24(2):180–190, 1928. doi: 10.1017/S0305004100015681.
- Christian Franzke. Predictability of extreme events in a nonlinear stochastic-dynamical model. *Phys. Rev. E*, 85:031134, Mar 2012. doi: 10.1103/PhysRevE.85.031134. URL <https://link.aps.org/doi/10.1103/PhysRevE.85.031134>.
- Christian L. E. Franzke. Extremes in dynamic-stochastic systems. *Chaos*, 27(1):012101, 2017. doi: 10.1063/1.4973541. URL <https://doi.org/10.1063/1.4973541>.
- Christian L. E. Franzke, T. O’Kane, J. Berner, P. Williams, and V. Lucarini. Stochastic climate theory and modelling. *WIREs Climate Change*, 6:63–78, 2015.
- Ana Cristina Moreira Freitas, Jorge Milhazes Freitas, and Mike Todd. Hitting time statistics and extreme value theory. *Probability Theory and Related Fields*, 147(3-4):675–710, 2010.
- Vera Melinda Gálfi, Tamás Bódai, and Valerio Lucarini. Convergence of Extreme Value Statistics in a Two-Layer Quasi-Geostrophic Atmospheric Model. *Complexity*, 2017:20 pages, 2017. doi: <https://doi.org/10.1155/2017/5340858>. URL <https://doi.org/10.1155/2017/5340858>.
- Eric Gilleland and Richard W. Katz. extRemes 2.0: An extreme value analysis package in R. *Journal of Statistical Software*, 72(8):1–39, 2016. doi: 10.18637/jss.v072.i08.
- B. Gnedenko. Sur la distribution limite du terme maximum d’une serie aleatoire. *Annals of Mathematics*, 44(3):423–453, 1943. ISSN 0003486X. URL <http://www.jstor.org/stable/1968974>.
- John Harlim. Model Error in Data Assimilation. In Christian L. E. Franzke and T. J. O’Kane, editors, *Nonlinear and Stochastic Climate Dynamics*, chapter 10, pages 276–317. Cambridge University Press, Cambridge, 2017.
- Klaus Hasselmann. Stochastic climate models part i. theory. *tellus*, 28(6):473–485, 1976.
- Mark P. Holland, Renato Vitolo, Pau Rabassa, Alef E. Sterk, and Henk W. Broer. Extreme value laws in dynamical systems under physical observables. *Physica D: Nonlinear Phenomena*, 241(5):497 – 513, 2012. ISSN 0167-2789. doi: <https://doi.org/10.1016/j.physd.2011.11.005>. URL <http://www.sciencedirect.com/science/article/pii/S0167278911003095>.
- Jonathan RM Hosking. L-moments: analysis and estimation of distributions using linear combinations of order statistics. *Journal of the royal statistical society. Series B (Methodological)*, pages 105–124, 1990.
- James L Kaplan and James A Yorke. Chaotic behavior of multidimensional difference equations. In *Functional Differential equations and approximation of fixed points*, pages 204–227. Springer, 1979.
- E.N. Lorenz. Predictability: a problem partly solved. In *Seminar on Predictability, 4-8 September 1995*, volume 1, pages 1–18, Shinfield Park, Reading, 1995. ECMWF, ECMWF.
- Valerio Lucarini, Davide Faranda, Jeroen Wouters, and Tobias Kuna. Towards a general theory of extremes for observables of chaotic dynamical systems. *Journal of Statistical Physics*, 154(3): 723–750, Feb 2014. ISSN 1572-9613. doi: 10.1007/s10955-013-0914-6. URL <https://doi.org/10.1007/s10955-013-0914-6>.
- Valerio Lucarini, Davide Faranda, Jorge Milhazes Freitas, Mark Holland, Tobias Kuna, Matthew Nicol, and Sandro Vaienti. *Extremes and recurrence in dynamical systems*. John Wiley & Sons, 2016.

- Hazime Mori, Hirokazu Fujisaka, and Hideto Shigematsu. A new expansion of the master equation. *Progress of theoretical physics*, 51(1):109–122, 1974.
- Arnold Neumaier and Tapio Schneider. Estimation of parameters and eigenmodes of multivariate autoregressive models. *ACM Transactions on Mathematical Software (TOMS)*, 27(1):27–57, 2001.
- T.N. Palmer and P.D. Williams. Introduction. Stochastic physics and climate modelling. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 366(1875):2419–2425, 2008. doi: <https://doi.org/10.1098/rsta.2008.0059>. doi: 10.1098/rsta.2008.0059.
- James Pickands. Statistical inference using extreme order statistics. *The Annals of Statistics*, 3(1):119–131, 1975. ISSN 00905364. URL <http://www.jstor.org/stable/2958083>.
- David Ruelle. Differentiation of SRB States. *Communications in Mathematical Physics*, 187(1): 227–241, Jul 1997. ISSN 1432-0916. doi: 10.1007/s002200050134. URL <https://doi.org/10.1007/s002200050134>.
- David Ruelle. A review of linear response theory for general differentiable dynamical systems. *Nonlinearity*, 22(4):855, 2009. URL <http://stacks.iop.org/0951-7715/22/i=4/a=009>.
- Gabriele Vissio and Valerio Lucarini. A proof of concept for scale-adaptive parametrizations: the case of the lorenz '96 model. *Quarterly Journal of the Royal Meteorological Society*, 144(710): 63–75, 2018. doi: 10.1002/qj.3184. URL <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.3184>.
- Daniel S. Wilks. Effects of stochastic parametrizations in the lorenz '96 system. *Quarterly Journal of the Royal Meteorological Society*, 131(606):389–407, 2005. doi: 10.1256/qj.04.03. URL <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1256/qj.04.03>.
- Jeroen Wouters and Valerio Lucarini. Disentangling multi-level systems: averaging, correlations and memory. *Journal of Statistical Mechanics: Theory and Experiment*, 2012(03):P03003, 2012. URL <http://stacks.iop.org/1742-5468/2012/i=03/a=P03003>.
- Jeroen Wouters and Valerio Lucarini. Multi-level dynamical systems: Connecting the ruelle response theory and the mori-zwanzig approach. *Journal of Statistical Physics*, 151(5):850–860, Jun 2013. ISSN 1572-9613. doi: 10.1007/s10955-013-0726-8. URL <https://doi.org/10.1007/s10955-013-0726-8>.
- Jeroen Wouters and Valerio Lucarini. *Parametrization of Cross-scale Interaction in Multiscale Systems*, chapter Chapter 4, pages 67–80. 2016. doi: 10.1142/9789814579933_0004.
- Robert Zwanzig. Ensemble method in the theory of irreversibility. *The Journal of Chemical Physics*, 33(5):1338–1341, 1960.
- Robert Zwanzig. Memory effects in irreversible thermodynamics. *Physical Review*, 124(4):983, 1961.