

On an extension of the L-moment approach to modelling distributions which include trend

D. A. Jones

ABSTRACT

In this paper, extensions of existing L-moment procedures are considered, aiming for ways of assessing changes over time in the distributions of quantities such as annual maxima while avoiding subdivision of data records into possibly overlapping sub-periods. Distributional changes of many types are included: location, scale, skewness, etc. Although direct application of some simple ideas is unsuccessful at providing good estimates of trends, less direct application of the same ideas allows for the implementation of significance tests for changes in shape, for the construction of confidence regions for the quantiles of a distribution that might change with time and for an indirect use of L-moments in estimating trends.

Key words | L-moments, trends

D. A. Jones
Centre for Ecology & Hydrology,
Maclean Building,
Benson Lane,
Crowmarsh Gifford,
Wallingford,
Oxfordshire,
OX10 8BB,
UK
E-mail: daj@ceh.ac.uk

BACKGROUND

The L-moment approach to analysing the statistical distributions of environmental variables has become widespread since its formulation by Hosking (1990) and its subsequent advocacy by Hosking & Wallis (1997). While an important use of this approach is as a means for fitting a selected family of statistical distributions to a set of data, it has found widespread use as a way of comparing and grouping the distributions appropriate to data at different sites, without having to make a choice of any one family of distributions at an early stage. L-moments provide convenient measures of the shape of distributions which can be compared numerically and graphically, and averaged across sites to yield an ‘average shape’.

When data are being assessed for possible trends, it would be convenient to have tools with similar properties to those of L-moments, whereby analyses could be made to examine changes in the shape of the statistical distribution of a variable over time, without having to make an initial assumption about the form of an appropriate family of distributions. There is a danger that a poorly chosen family of distributions might lead to poor conclusions, or that important features of the behaviour of data might be missed. Thus a distribution-free approach to assessing trends would

be an important adjunct, even if a final analysis was eventually made using a chosen family of distributions.

The existing methodology of L-moments is based on the assumption that data arise from a fixed distribution and thus it is not immediately suited to assessing trends in the shapes of distributions. Chandler & Scott (2011) provide a general survey of statistical techniques for detecting trends and they take the view, as is done here, that ‘trend’ refers to changes in all aspects of statistical distributions, not just a location parameter. The purpose of this paper is to consider how the L-moment methodology can be modified to allow the analysis of trends, without breaking datasets into blocks (as has been done in some applications; see below), such that the analysis reduces to the existing L-moment method if no trends are fitted and such that the computational simplicity of the L-moment approach is retained as far as possible. Foundation of the approach on L-moments means that any methods developed would be targeted to those qualities of the distributions being analysed that are quantified by the L-moments, or by the L-moment ratios, which include versions of location, scale, relative scale, skewness and kurtosis.

The following instances are examples of some previous uses of L-moments in connection with trend analyses.

doi: 10.2166/nh.2012.081

Adamowski & Bougadis (2003) used L-moments to delineate homogeneous collections of sites for extreme rainfalls (of various durations) in Canada, before undertaking trend analyses that did not involve the L-moment approach and that did not consider trends in the shape of the distributions. Fowler & Kilsby (2003) used L-moments for extreme multi-day rainfalls in the UK, first to form groups of sites and then as part of a trend analysis that involved obtaining L-moment estimates from data both within moving 10-year windows, and within fixed decadal windows. They investigated trends firstly by plotting these estimated L-moments against time, and then by showing changes in the rainfall-frequency curves fitted using the L-moment estimates within each decade. Kay & Jones (2012) analysed changes to the distribution of annual maximum river flow derived from rainfall-runoff models driven by the rainfall output at a 1 hour resolution from a transient climate change model; they used L-moments estimated from 30-year moving windows and then fitted non-parametric trends to these and to the quantiles of distributions fitted within each 30-year window. They went on to use a resampling technique to assess the statistical significance of the apparent changes.

One objective of the present work is to avoid the apparent necessity of using blocks of data within the overall dataset in order to examine trends in L-moments. A very naïve method of fitting trends using L-moments is proposed that, while it has extremely poor properties regarding quantifying changes over time in the shapes of distributions, does yield good estimates when no changes are present. This important finding leads to the development of two useful types of analysis: (i) tests for presence of changes in distribution; and (ii) confidence regions for the quantiles of distributions that possibly change over time.

BASIS OF TREND ESTIMATION WITH L-MOMENTS

In order to avoid unnecessary formulae when these are obvious, it will be convenient to use the third L-moment, or later the L-skewness, as the basis of discussion, on the understanding that other cases can be being dealt with in a similar way. A route to applying a regression-like analysis of trends to L-moments can be developed using the formulation of the usual sample estimate of an L-moment as

U-statistic, which is noted by Hosking (1990; Section 3.1). Thus, consider the estimation of the third L-moment, λ_3 , defined in the standard notation by:

$$\lambda_3 = \frac{1}{5}E(X_{3:3} - 2X_{2:3} + X_{1:3}) \quad (1)$$

where it supposed that three independent samples of a random variable, X , are obtained and put into decreasing order as $X_{3:3}, X_{2:3}, X_{1:3}$. In this often used notation, $X_{j:k}$ denotes the j 'th largest value in a set of k values of a quantity originally designated with the symbol 'X'. For a given dataset, $\{x_j\}$, of n items, a new set of values, $\{d_{ijk}\}$, is created, where the indices i, j and k satisfy $1 \leq i < j < k \leq n$, and where the items d_{ijk} are defined as follows. For the given values of i, j and k , the values (x_i, x_j, x_k) are put into decreasing order and labelled temporarily as (y_3, y_2, y_1) : then d_{ijk} is defined as:

$$d_{ijk} = \frac{1}{5}(y_3 - 2y_2 + y_1) \quad (2)$$

The set $\{d_{ijk}\}$ contains N items, where in this case $N = \frac{1}{6}n(n-1)(n-2)$. (The value of N varies with the order of the L-moment being considered and, for example, $N = n$ for the first-order moment.) The standard sample third-order L-moment estimator, l_3 , can be then written as:

$$l_3 = N^{-1} \sum_{i < j < k} d_{ijk} \quad (3)$$

The quantities d_{ijk} can conveniently be called the third-order raw L-moments. Other orders of L-moments can be treated similarly by treating ordered subsets with 1, 2, 3, ... members (equivalent to the number of subscripts on the raw L-moments, d) for the L-moments of order 1, 2, 3, The first-order raw L-moments are identical to the original data values.

The relationship of the ordinary sample L-moment estimates to regression can be seen by noting that the sample average in Equation (3) would be the result of the estimate for a regression coefficient obtained by ordinary least squares for a model in which the N values d_{ijk} are the dependent variables and in which there is a single independent variable with constant values, with value equal to 1. That

is, the estimate l_3 is the value of u which minimises the sum of squares:

$$S(u) = \sum_{i,j,k} \{d_{ijk} - u\}^2 \quad (4)$$

Here, the ordinary least-squares step would be employed simply as a means of minimising a sum of squares. The usual assumptions about uncorrelatedness or independence of error terms that are often associated with least-squares procedures do not hold here, but they are not needed as none of the results deriving from them is required. In order to apply a regression-like analysis for trend in L-moments, it is required to have a set of values μ_{ijk} that would reflect what would be the expected outcomes of d_{ijk} under a given model of trend. For the least-squares estimation to be simplest, these values μ_{ijk} would be expressed as a linear combination of fixed regression coefficients with dependent variables that could vary with all of the subscripts i , j and k . The structure of the model here would be derived from an assumed model for trend in the L-moments.

The raw L-moments for a set of data that is used again later are shown in Figure 1. Here the first three orders of L-moments are shown and the raw L-moments are plotted against a notional time-point equal to the average of the times from which the underlying values derive. The dataset consists of 70 years of annual maximum flow values at Day's weir on the River Thames, and Figure 1(a) shows these values as the first-order raw L-moments.

Someone who is considering the possibility of a change in the shape of a statistical distribution, within the context of L-moments, would most naturally start from the possibility that the L-moments, such as the L-moment in Equation (1), might vary with time, and might start with a regression-like equation expressing this formally, for example as:

$$\lambda_3(s) = \theta_0 + \theta_1 z_s \quad (5)$$

Here θ_0 and θ_1 would be regression coefficients to be estimated and z_s represents a known function of the time, s , such as a linear function or a step-function, which would specify the type of trend being contemplated. However, while it is possible to define the quantities μ_{ijk} using a version of Equation (1) in which the three underlying values of X are

drawn from three different distributions, there is no direct formula linking them to the third L-moments of the separate distributions, and hence there is no expression for them starting from Equation (5). Specifically, a formal expression for μ_{ijk} requires complicated integrals of the three different distribution functions of the component random variables which, for general distributions, cannot usefully be represented in terms of the sets of L-moments of the three distributions. By assuming that the distributions are in some particular family of distributions, values for μ_{ijk} could eventually be obtained from the L-moments of the separate distributions, but even then they would typically need to be obtained by numerical integration. As indicated in the introduction, the intention here is to investigate computationally simple approaches to trend analysis, which rules out this 'exact' approach. There is also a wish to avoid the assumption that a particular family of distributions holds at too early a stage in the process. Instead, simple approximations are pursued in the hope that the effects of the errors introduced will be small. As indicated later, this turns out not to be the case. The simple approximations here can be criticised on the basis that they ignore any changes in the location and scale parameters but no simple formula involving these is known. Nonetheless, the approach here does still provide a basis for useable methods for statistical inferences of many types, and these are described later.

The type of approximation made is that the cross-time L-moments, such as μ_{ijk} , can be replaced by simple averages of the corresponding single-time L-moments $\lambda_3(t)$ at the corresponding times i , j and k . Then the regression model in Equation (5) for the single-time L-moment would result in:

$$\mu_{ijk} \cong \theta_0 + \theta_1 \frac{1}{3}(z_i + z_j + z_k) \quad (6)$$

and, given this, it would then be possible to estimate θ_0 and θ_1 using standard least-squares fitting routine with the set of values d_{ijk} as the dependent variables and with independent variables consisting of a constant set of values set at 1, and a set of values equal to $\frac{1}{3}(z_i + z_j + z_k)$. Clearly the same approach can be taken for generalised versions of Equation (5) having more than one explanatory variable z , and hence more regression coefficients.

This approach needs to be modified slightly to allow the trend analysis using L-moments to follow the established




Figure 1 | Raw L-moments of orders 1, 2 and 3, plotted at average time-points. Data are annual maxima for Day's weir 1938–2007.

practice of using L-moment ratios, rather than L-moments themselves, to compare distributions.

TREND ANALYSIS FOR L-MOMENT RATIOS

Data analyses that involve L-moments typically make use of L-moment ratios rather than the L-moments themselves, because these dimensionless measures seem more useful for comparison of distributions at different sites. In the case of trend analysis it may be most natural to consider trends in shape parameters, such as the L-skewness, but dealing with the scale parameter is less obvious. There are two simple possibilities in the case of trends in annual maximum flood: firstly, that a change in the mean level is a reflection of a simple upward shift of the distribution, so that the L-scale would remain fixed; secondly, that the distribution is affected in a proportional manner, so that the coefficient of L-variation (L-CV) would remain fixed. Thus trend analyses of either or both of L-scale and L-CV may be appropriate in any particular case.

It is convenient to label the quantities λ_1 , τ , τ_3 , etc. as the 'L-moment ratios' even though λ_1 , which is the mean, is not formally an L-moment ratio. Here τ is the L-CV and τ_3 is the L-skewness in the notation of Hosking & Wallis (1997). Once again, the discussion here is primarily in terms of the third-order quantity. It is assumed that models have been fitted to the lower-order L-moments (or L-moment ratios) and that these models are now treated as fixed. In order to determine what type of regression-fitting is appropriate for L-moment ratios, use can be made of the requirement that the result of the method when applied in a trend-free case should reduce to the standard formula for an L-moment ratio. In the case of the L-skewness, the standard estimate is:

$$t_3 = l_3/l_2 \quad (7)$$

where l_2 is the sample estimate of the L-scale and l_3 is the sample estimate of the third L-moment. It follows that the usual estimate for the τ_3 is identical to the value of u which minimises the sum of squared errors defined as:

$$S(u) = \sum_{i,j,k} \{d_{ijk} - l_2 u\}^2 \quad (8)$$

Given this situation, a trend in the L-skewness would be estimated by constructing values μ_{ijk} based on the existing model for l_2 (as a function of time) and the model to be fitted for the L-skewness. Here μ_{ijk} have exactly the same interpretation as earlier, as the values that d_{ijk} estimates. In parallel to Equation (5), suppose that the model for the L-skewness at time s is of the form:

$$\tau_3(s) = \theta_0 + \theta_1 z_s \quad (9)$$

and that l_2 is a known function of time, h_s . Then two ways of approximating μ_{ijk} based on the idea of using average values of the parameters of the distributions at different times are:

$$\begin{aligned} \mu_{ijk} \cong & \frac{1}{3}(h_i + h_j + h_k) \theta_0 \\ & + \frac{1}{3}(h_i + h_j + h_k) \frac{1}{3}(z_i + z_j + z_k) \theta_1 \end{aligned} \quad (10)$$

$$\mu_{ijk} \cong \frac{1}{3}(h_i + h_j + h_k) \theta_0 + \frac{1}{3}(h_i z_i + h_j z_j + h_k z_k) \theta_1 \quad (11)$$

where the first corresponds to using averaged values of λ_2 and τ_3 to determine the approximating distribution and the second corresponds to using average values of λ_2 and λ_3 . In either case, estimation of θ_0 and θ_1 can again be achieved using a standard least-squares fitting routine. This would have d_{ijk} as the dependent variables and would have independent variables consisting of the values multiplying θ_0 and θ_1 in Equations (10) or (11). A similar third expression for μ_{ijk} can be found that would correspond to using averaged values of λ_1 , τ and τ_3 .

PRACTICAL CONSIDERATIONS

While the fitting of models for trends in L-moments or in L-moment ratios is computationally simple using the approach outlined above, there are various considerations needed in making use of the results. These have been noted in a trial implementation of the approach. Some of these considerations relate to the bounds that need to be satisfied by values of L-moments and their ratios so that they are possible values pertaining to a probability distribution. Another finding relates to the lack of resilience of the estimated trends to outliers at the stage of fitting by least squares. A further

consideration relates to whether a simple model for trend in the L-moments behaves in a reasonable way when interpreted via the quantiles of a given family of distributions.

Hosking & Wallis (1997; Section 2.6) outline a number of bounds that must be obeyed by L-moments and their ratios. There is no reason why any of these bounds should be automatically obeyed by the results of fitting a trend model by the regression-like procedures discussed here. The simplest example arises when considering the model for the first L-moment (equivalent to the mean) when a simple linear-in-time trend is fitted: the procedure here is identical to the well known procedure of fitting a linear trend in typical values by least squares. Even if the population from which observations arise contains only positive values, as in the case of annual flood maxima, the fitted linear trend will produce negative values at some point in time except in the trivial case. Similarly, if trends in λ_1 and τ are fitted separately, then the values of λ_2 derived via the relation: $\lambda_2(s) = \lambda_1(s)\tau(s)$ may not satisfy the bound $\lambda_2(s) \geq 0$ that must hold if $\lambda_2(s)$ is to be treated as the L-scale of a probability distribution; values contradicting the bound may arise from negative fitted values of either of $\lambda_1(s)$ or $\tau(s)$. Similarly to λ_1 , fitting a trend model for λ_2 may also yield negative values in the fitted trend function, so that neither L-scale or L-CV has an advantage in this regard. There are other relevant bounds on the ranges of L-moment ratios that must be satisfied for the fitted trends in L-moments to correspond to a valid probability distribution. One of these bounds is $|\tau_3| < 1$, and this must apply either for a trend model fitted for τ_3 , or for the ratio of values derived from separate models for λ_2 and λ_3 .

Clearly, if the fitted trends in L-moment ratios are to be taken forward for use in producing corresponding trends in quantiles via a selected family of distributions, the theoretical bounds will need to be obeyed. The work here has taken the most immediately convenient approach in implementing a trend-fitting procedure. Specifically, an initial version of the trend function is fitted by linear least squares, as outlined above, and then a final version of the fitted trend at any given time is created by trimming the values from this fitted function, so that the bound is obeyed. When a fitted trend in a lower-order L-moment is required for use as fixed values in fitting a model for a higher-order L-moment ratio, it is these trimmed values that are carried forward.

Similarly, it is the trimmed values that are used as the values for fitting a distribution for any given time-point. Thus, even if models for changes in the L-moments that are linear in time are chosen for fitting, the final results would relate to the trimmed versions of these models.

An alternative to the above convenient approach would be to formulate trend models that are guaranteed to yield values within the valid ranges, but this goal cannot be achieved within the linear regression framework being used. However, the methodology could be modified by allowing the use of readily available computer procedures for non-linear least-squares estimation, at the expense of a modest increase in computation time. Such an approach would also provide a way of including the trimmed model discussed more integrally within the fitting procedure, compared to the 'convenient' method. However, it would usually be more natural to adopt a smooth function of time as a model for trend.

For applications to annual flood maxima, the results from a standard least-squares approach to fitting a trend in the mean value are likely to be affected by outliers simply because the probability distributions of such flood values are typically long tailed. The same has been found to be true in the fitting of trends to the higher-order L-moments by the approach described here. For example, a single outlier in a time series of n items will affect $\frac{1}{2}(n-1)(n-2)$ out of the $N = \frac{1}{6}n(n-1)(n-2)$ values of d_{ijk} that would be used in the least-squares fitting for the third L-moment. This finding would suggest that an outlier-resistant approach to fitting the trend models should be adopted for fitting each of the L-moments. If one maintains the requirement that the result of fitting a model with no trend should produce results identical to those from the standard L-moment estimates, this would mean that resistant-fitting procedures would need to be specifically selected. Existing outlier-resistant regression procedures would affect all of the regression coefficients, including the constant term, and it is this term that corresponds to the standard L-moment estimates.

A final practical consideration is whether simple trend modelling of the shapes of distributions using L-moments does correspond to the distributions behaving in a way that a data analyst would expect. In particular, a monotonic trend in the L-skewness does not imply a corresponding monotonic behaviour in the quantiles of a fitted distribution. For example, suppose that the model being

considered is such that the mean and L-scale are constant, but that the L-skewness has a trend. Suppose also that a generalised logistic distribution is to be applied at some later stage; similar results arise for other distributions. Figure 2 shows how the quantiles of this distribution vary as a function of the L-skewness, for $\lambda_1 = 1$ and $\lambda_2 = \frac{1}{2}$, where each line represents the values that would be plotted on the Gumbel scale at a given value of the 'reduced variate', y , where $y = -1, 0, 1, \dots$. This behaviour arises because, within this family (and similar families) of distributions, for fixed mean and L-scale, an increase in the skewness means that a greater proportion of the distribution is brought close to the centre while the remaining part of the distribution is pushed even further way from the centre. Eventually, as the skewness increases, the quantile for any fixed return period becomes close to the median. These types of behaviour mean that modelling trends using L-moments might not actually have the interpretation that is expected. Mathematically, the behaviour arises because the scale parameter of the generalised logistic distribution depends on both L-scale and L-skewness.

ESTIMATION OF TRENDS

A limited simulation experiment of the properties of the trend estimation scheme described above has been carried




Figure 2 | Anomalous behaviour of the quantiles of a distribution. Showing quantiles of a generalised logistic distribution with fixed first and second L-moments as the L-skewness varies. Each line is for a different return period and is labelled with the corresponding Gumbel reduced variate.

out, of which only a summary is reported here. The experiment involved generation of independent random samples of length 201 from generalised logistic distributions in which the time-varying parameters were determined from models for the L-moment ratios (the mean, L-CV and L-skewness), and these same models were then fitted to the data values generated. The models selected allowed each of the three L-moment ratios to vary as a linear function of time, with no connection between the parameters of these models. Estimation of these trend models followed the scheme described earlier, and included the trimming of the estimated trends that was outlined. Two versions of the estimation scheme were outlined, as exemplified by Equations (10) and (11), and both have been implemented. The notional record length of 201 is, of course, rather longer than records of annual maximum flood typically available in hydrology, but it represents a choice that gives the methodology for estimating trends in shape a good chance of working. Choosing an odd number makes it easy to examine the properties of estimates at the mid-point of the series.

The results found are described only briefly here. While the estimation schemes associated with Equations (10) and (11) did produce different results, neither was consistently better than the other, and the difference between them was found to be considerably smaller than their inherent biases, which are discussed later. The results for the trend in the mean are as would be expected from theory, that the regression coefficients for the mean are unbiased for any amount of trend in the L-CV or L-skewness.

Results if there is no trend present

If there is no trend in any of the three L-moment ratios in the model used for generating the data series, then the estimated parameters in the trend models perform well, being nearly unbiased. Analysis of 40,000 repetitions of the simulation procedure demonstrates that bias is definitely present in the estimates of some of the coefficients, but this bias is small compared with the standard error of an individual estimate (less than 5%). Similar findings apply to other quantities that might be studied but which vary with time; the fitted L-moments, L-moment ratios and quantiles of fitted generalised logistic distributions, with the worst bias reaching 10% of the estimate's standard error. The biases in the

estimates of the regression coefficients of the L-CV and the L-skewness, which are only the 'constant' terms, arises because of the use of estimated values from an earlier stage in the overall estimation procedure. The reason for the bias parallels that for the bias in the usual estimates of L-moment ratios: for example, in Equation (7), the estimates l_2 and l_3 are unbiased, but the result, t_3 , is a biased estimate of τ_3 .

Results if trend is present

If there is trend in one or more of the L-moment ratios, then more substantial biases exist, with the amount of bias depending on the size of the trends. A trend in the mean leads to biases in all the estimates of the regression coefficients of the L-CV and L-skewness. A trend in L-CV leads to biases in the estimates of the 'constant' regression coefficient for both L-CV and L-skewness. However, a trend in L-skewness does not itself lead to biases in any of the regression coefficients. Here an estimate is being called biased if the bias is greater than 10% of the estimate's standard error. If the trends are large enough to be seen visually in a time-series plot, then the biases can exceed three times the standard error. This applies not only to the estimates of the regression coefficients but also to the estimates of the time-varying quantiles derived by assuming that the generalised logistic distribution is known to apply.

Conclusion

Estimation of trends in the shapes of distributions using the methodology outlined here cannot be recommended. The estimates might be of use in providing initial values for some more sophisticated procedure in which a particular family of distributions is selected at an early stage, for example using maximum likelihood estimation.

One obvious reason for the method's poor performance as an estimation procedure in the presence of trend is the approximation made in Equations (6), (10) and (11). The most immediate remedies have been ruled out by the requirement to keep the procedure simple and assumption-free. However, useable improved approximations might yet be developed. The differences between the results for schemes based on the approximations in Equations (10)

and (11) are very small, and it is judged unlikely that any variant of these kinds will produce an improvement.

TESTING FOR TREND

In spite of the poor results reported above regarding the performance of the methodology as an estimation procedure, there is still a possibility of using the approach to provide a test for trend-like changes in the shapes of distributions. This possibility derives from the result that the regression estimates are effectively unbiased if no trends are present. Thus a test statistic can be constructed based on how far the method's estimated trend is from 'no trend'; the method produces results that indicate more trend if more trend is present, even if its quantification of the amount of trend present does not agree with the amount actually present. It is envisaged that a permutation test will be suitable for analysing annual flood maxima where the assumption of statistical independence between years is often made. If independence cannot be assumed then a block-bootstrap procedure would provide an alternative. Permutation procedures are well known in statistics (Maritz 1981; Good 1994). In hydrology, Robson *et al.* (1998) and Robson & Reed (1999; Section 21.5) applied the method to testing for trends in the typical sizes of quantities derived from flood data.

To provide a single overall test for change across time in the distributions being analysed, where there might be any type of change in these distributions, the initial requirement is to determine a single test statistic derivable from an observed data series, to which a permutation test can be applied. This test statistic should be something that quantifies the apparent amount of change in the distributions that an analysis of the data series finds. If the estimation methodology described here is applied for L-moments up to the third order, then the three stages of estimation each might contribute a measure of change, but there is no immediate way of combining them into a single measure of change. Thus the question is how to combine a measure of change in mean level, a measure of change in L-CV and a measure of change in L-skewness. Here the basic measures might be either trend coefficients themselves or the reduction in the sums of squares achieved by fitting each trend model. The approach suggested here is to make a

dual use of the permutation samples. For this method, the permutations are used to derive a probability of getting a more extreme value for each of three tests statistics, one for each L-moment ratio. These three values are then combined by taking the smallest, to produce the combined test statistic, and finally the permutation approach is applied again to evaluate the true probability of getting a smaller result for the smallest of the three values.

For the dataset of annual maximum floods at Day's weir (Figure 1), estimates of the trend coefficients for simple linear trends in λ_1 , τ and τ_3 were all negative. The permutation analysis outlined above indicated that the absolute values of the trend coefficients such that the probabilities of getting more extreme values than those observed were 39.9, 5.5 and 11.9%; this is according to their permutation distributions representing no trend. Judged individually these results might be interpreted as pointing to a conclusion that there is a downward trend in L-CV, with no trend in mean value for the annual maxima and a weak indication of a reduction in skewness. However, the second stage of the permutation analysis indicates that in fact there is a 15.5% chance of getting an even smaller value than the observed value of 5.5% for the smallest of these three probabilities (i.e. the combined test statistic) if there is no trend. Thus, if one starts from the position of not knowing what type of change in the distribution of annual maxima might have occurred, the test does not provide clear evidence of any change: the test statistic would be just significant at a significance level of 15.5%.

As described earlier, the trend model fitted at any stage can be badly affected by outliers and how the raw fitted trend may produce invalid values for L-moments or L-moment ratios. This raises some doubt about using the estimated trend coefficients as a measure of the amount of change. One proposal would be to use the trimmed versions of the fitted time-varying functions as the final estimates, just as they are within the procedure to carry forward to the next stage. The amount of change detected at a given stage could then be quantified by calculating the sum of squares of deviations of this function about its own mean. A similarly constructed test statistic was used in the test of trend employed by Kay & Jones (2012). The same idea may be of use if the model for trend included more than just a single parameter controlling departure from a constant value.

The approach to testing for trends based on L-moment ratios is moderately complicated in a computational sense, as the analysis of the higher-order L-moment ratios for each permutation requires a regression analysis with a new set of independent variables, as these are constructed as a result of the previous stage of the analysis. Some of this complexity can be avoided if an analysis is made of the L-moments directly. However, a test of trend based on simple models for the L-moments would be most sensitive to a different pattern of possible trend than would tests using simple trend models for the L-moment ratios. For comparison with the results of the test for linear trends in λ_1 , τ and τ_3 in the data for Day's weir, the test for linear trends in λ_1 , λ_2 and λ_3 gave exceedance probabilities of 39.9, 3.7, and 11.5% for the individual components, a combined test statistic of 3.7% and this number was just significant at a significance level of 10.7%.

To summarise, the results for tests for changes in L-moment ratios or in unscaled L-moments are broadly similar in that they both indicate a change in scale, rather than location or skewness as the predominant effect. Once account is taken of having done tests of change in several different aspects of the distribution of annual maxima, the overall test statistics are not significant enough to draw a definite conclusion that change has occurred, but close enough to significance to prompt further analyses of how much change is represented within the dataset.

CONFIDENCE INTERVALS FOR TREND

A simple significance test for change in shape of distributions may not be adequate for supplying information about the amounts of change that might be occurring in the shapes of the distributions underlying a particular set of observations. Visual assessment of the amount of change apparent in a dataset can perhaps be most readily provided in the form of confidence intervals on the quantiles in the distributions, where these are allowed to vary in time.

The approach suggested here is based on using a selected family of distributions, in this case the generalised logistic distribution, and on representing the possible changes in shape by allowing linear-in-time changes to the parameters of this distribution (using the usual

parameterisation). While this approach clearly does make the type of distributional assumption that was earlier being avoided, it does make use of the L-moment-based approach to assessing trends that is described in this paper. The approach has been chosen on the basis of being achievable with moderate computational resources. At the centre of the approach is a significance test of the hypothesis that a series of data is a trend-free realisation from a standard generalised logistic distribution; specifically one where the distribution's location, scale and shape parameters are zero, one and zero. A confidence region for a set of parameters, including trend parameters, is then constructed by including within the confidence region any such set of parameters for which there is acceptance of the significance test just mentioned when applied to a series of values obtained from the observed data series by the following de-trending and standardisation step. Values y_t in the series to which the test is applied are calculated from values in the observed series, x_t , as:

$$y_t = F^{-1}(F(x_t; u_t, \alpha_t, k_t); 0, 1, 0) \quad (12)$$

where $F(x; u, \alpha, k)$ is the cumulative distribution function of the generalised logistic distribution with parameters u , α and k , with $F^{-1}(x; u, \alpha, k)$ being the inverse of this function. The values u_t , α_t and k_t represent time-varying values of the parameters derived from the trend model. If the trend model is correct, then the values calculated as $F(x_t; u_t, \alpha_t, k_t)$ would correspond to a uniform distribution on the interval 0 to 1, and it would be possible to base the confidence intervals on tests of this distribution, rather than including the further step of transforming this distribution to a standard generalised logistic distribution. This alternative approach has not been tried and the one adopted has been chosen on the dual basis that L-moments are widely accepted for use with extreme value like distributions, but not for use with the uniform distribution, and of leaving the skewness of the distribution nearly unchanged if the values of the shape parameter are close to zero. The above expression simplifies to:

$$y_t = \begin{cases} -k_t^{-1} \log\{1 - k_t \alpha_t^{-1} (x_t - u_t)\}, & k_t \neq 0, \\ \alpha_t^{-1} (x_t - u_t), & k_t = 0. \end{cases} \quad (13)$$

It may be noted that, if the distribution here were replaced by a generalised extreme value distribution, the same transformation would apply.

The test that the constructed series represents a trend-free realisation from a standard generalised logistic distribution is based on a similar idea to that used above, but modified to test for specific values of the L-moments, as well as for 'no trend' in the L-moments. In this application, in which the constructed series are meant to have a mean value close to zero, it seems appropriate to test for trends in the L-moments, rather than in the L-moment ratios, and no trimming is applied to force the estimates to jointly satisfy theoretical bounds. The method previously outlined is applied, with a linear-in-time model being fitted for each of the first three L-moments, yielding a total of six regression parameters. To be specific, the regression models here are formulated in such a way that, for example in Equation (5), the 'constant' regression parameter θ_0 represents the average of the values of $\lambda_3(s)$ across the set of time-points being analysed. The values that the six regression parameters should take are determined by the hypothesis being tested: they should all be zero except for the 'constant' regression parameter for L-scale model, which should be one.

The permutation-based resampling approach to constructing the test procedure used to test for trend is here replaced by straightforward simulations from the standard generalised logistic distribution on which the model for possible trends is based. Only a single set of simulations is required to determine critical values for the test here, with the simulations being done for the standard logistic distribution. The same critical values are applied to whatever values of u_t , α_t and k_t are being tested to see if they represent possible true values of the parameters of the time-varying distributions of observations at different times.

As for the test for trend, the question of combining information from the six test statistics for the individual regression coefficients into a final test statistic might be resolved by initially determining a probability of getting a larger departure from a true value for each of the six, and then finding what value of the smallest of these six probabilities will occur in a given percentage of samples of the same length as the observed series. However, an initial simulation study showed low correlations between the errors in these estimated parameters: they are all very close to zero apart from

correlations close to 0.3 for the two pairs of estimates: (i) the ‘constant’ regression parameters of the models for trends for λ_1 and λ_2 , and (ii) the trend coefficients for λ_1 and λ_3 . Given this situation, it is reasonable to adopt the simplified procedure of using a test statistic defined as being a weighted sum of squares of the differences between the estimated parameters and their true values under the null model. The initial simulations were used to determine the variances of the estimated parameters and these were used to specify weights inversely proportional to these variances in the weighted sum of squares. Such weights need not be determined particularly accurately, but they are kept fixed for later stages of the analysis. If the various errors were normally distributed (and if all the correlations were zero) the sum of squares would have a χ^2 distribution with six degrees of freedom, but the later simulations show that the test statistic has a longer tail than this. The validity of the distribution obtained from the simulations is not affected by either the non-normality of the estimates or the correlation that is actually present.

Linear-in-time models for the distribution parameters u_t , α_t and k_t are formulated, giving a total of six parameters controlling the variation of the distribution over time. The first step in constructing a confidence interval for the quantiles is to construct a collection of sets of the six parameter values such that these effectively determine the boundary of the six-dimensional confidence set for the parameters. Specifically, for a 90% confidence region, a search is made for values of the six parameters such that the de-trended and standardised series is just on the border of having the underlying significance test accepted or rejected at a 10% level. This collection of boundary values in six-dimensional space is then used to determine a corresponding collection of quantile-time functions that correspond to these ‘just acceptable’ parameter sets.

The approach here also allows a set of ‘central estimates’ for the model parameters. In this specific case this results in values such that the sample L-moments of the transformed values derived from Equation (13), but in other cases (assuming a weighted sum-of-squares test criterion is used), the central estimates could be defined as values minimising the test criterion. These central estimates are different from those considered in the section on ‘estimation of trends’.

Some results from this procedure are shown in Figure 3, which presents 90% confidence regions for flows having a

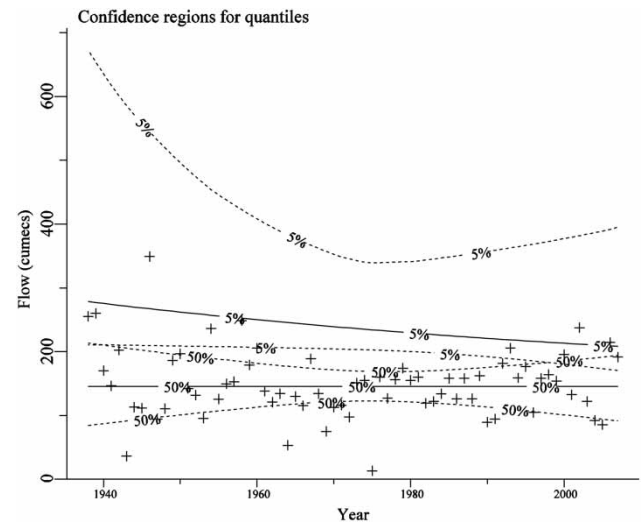


Figure 3 | Central estimates and 90% confidence regions (dashed lines) for flows having 50 and 5% exceedance probabilities in the given year. Showing annual maxima for Day’s weir 1938–2007.

50% and a 5% chance of being exceeded in each particular year for the 1938–2007 water-years for the Day’s weir dataset. Here each region should be interpreted as being the outcome of a procedure that had a 90% chance of the constructed region covering the true quantile curve for all years. Figure 4 has a similar interpretation, but shows a 90% confidence region for the ratios defined as the flow of a given return period in the last year of record (2007) divided by the flow of the same return period in the first year of




Figure 4 | Central estimate and 90% confidence region (dashed lines) for the ratio defined by the flow of a given return period in 2007 divided by the flow of the same return period in 1938.

record (1938). The reason for the narrowing of the confidence region near a return period of 10 years is unknown and is still the subject of speculation.

Further extensions

It is not immediately possible to use the calculations described above to derive confidence intervals for individual parameters or for single derived quantities such as the flow of a given return period in a particular year. The confidence levels quoted relate to whether the whole of the function being studied lies entirely within the confidence region. However, it should be possible to derive such 'pointwise' confidence intervals by finding test statistics for use within the procedure that are appropriate to this purpose.

As indicated above, the approach suggested here for constructing confidence intervals can be used as an indirect procedure for parameter estimation. In fact this has wider applications in situations not involving trends. It provides an alternative to the standard L-moments procedure, but based on the indirect use of L-moments, and will produce different results. The introduction of the step in Equation (13), which has a re-shaping effect as well as a de-trending effect on the distributions, makes it possible to consider fitting distributions for which there are no explicit formulae for the L-moments: such distributions would effectively be defined by the re-shaping function being used, on which there is little constraint. Such revised estimation procedures would need careful study.

A further possibility is to extend the procedures outlined here to develop tests of whether a given distribution is acceptable as a model for a given dataset. This development would give new types of L-moment-based tests of fit.

CONCLUSIONS

Simple basic ways of extending L-moment estimation procedures to assess trends in the shapes of distributions have been outlined. While a straightforward implementation of these procedures is not successful in direct estimation of such trends, they can be used in conjunction with

permutation techniques to provide significance tests of trend in distributional shape, and they can be used in conjunction with de-trending and standardisation steps to provide the basis for confidence regions for how quantiles of flow might change over time, and for related quantities. The approach opens new possibilities for estimating parameters. The extended procedures can be applied to estimation problems in which the models are such that data can be transformed to have a simple standard distribution, even if the L-moments of the untransformed data cannot be evaluated analytically. In addition, the use of parameterised transformations to represent trends in the shape of the distribution over time may have advantages compared to using the ordinary parameters of a distribution, as it allows more control over how changes in the quantiles of the distribution behave over time.

REFERENCES

- Adamowski, K. & Bougadis, J. 2003 [Detection of trends in annual extreme rainfall](#). *Hydrological Processes* **17**, 3547–3560.
- Chandler, R. E. & Scott, M. E. 2011 *Statistical Methods for Trend Detection and Analysis in the Environmental Sciences*. Wiley, New York.
- Fowler, H. J. & Kilsby, C. G. 2003 [A regional frequency analysis of United Kingdom extreme rainfall from 1961 to 2000](#). *International Journal of Climatology* **23**, 1313–1334.
- Good, P. 1994 *Permutation Tests: A Practical Guide to Resampling Methods for Testing Hypotheses*. Springer-Verlag, New York.
- Hosking, J. R. M. 1990 L-moments: analysis and estimation of distributions using linear combinations of order statistics. *Journal of the Royal Statistical Society, Series B* **52**, 105–124.
- Hosking, J. R. M. & Wallis, J. R. 1997 *Regional Frequency Analysis: An Approach Based on L-moments*. Cambridge University Press, Cambridge.
- Kay, A. L. & Jones, D. A. 2012 [Transient changes in flood frequency and timing in Britain under potential projections of climate change](#). *International Journal of Climatology* **32**, 489–502.
- Maritz, J. S. 1981 *Distribution-Free Statistical Methods*. Chapman & Hall, London.
- Robson, A. & Reed, D. 1999 *Flood Estimation Handbook, Volume 3: Statistical Procedures for Flood Estimation*. Institute of Hydrology, Wallingford.
- Robson, A. J., Jones, T. K., Reed, D. W. & Bayliss, A. C. 1998 [A study of national trend and variation in UK floods](#). *International Journal of Climatology* **18**, 165–182.