

L-MOMENT STATISTICS

3-1 L-MOMENTS

L-moment statistics are used extensively in L-RAP. L-moment statistics are used for computing sample statistics for data at individual sites; for testing for homogeneity/heterogeneity of proposed groupings of sites (regions); for conducting goodness-of-fit tests for identifying a suitable probability distribution(s); and for solving for distribution parameters for the selected probability distribution. In particular, estimation of the L-moment ratios of L-Cv and L-Skewness is a key element in determining the success of the regional frequency analysis in computing quantile estimates for selected sites.

L-moments obtain their name from their construction as linear combinations of order statistics (Hosking and Wallis¹⁴, pp18-27). They are a dramatic improvement over conventional product moment statistics for characterizing the shape of a probability distribution and estimating the distribution parameters, particularly for environmental data where sample sizes are commonly small. Unlike product moments, the sampling properties for L-moments statistics are nearly unbiased, even in small samples, and are near Normally distributed. These properties make them well suited for characterizing environmental data that commonly exhibit moderate to high skewness.

The L-moment measure of location, and L-moment ratio measures of scale, skewness and kurtosis are:

Location, mean:

$$\text{Mean} = L_1 \tag{3-1}$$

Scale, L-Cv (t_2):

$$t_2 = L_2/L_1 \tag{3-2}$$

L-Skewness (t_3):

$$t_3 = L_3/L_2 \tag{3-3}$$

L-Kurtosis (t_4):

$$t_4 = L_4/L_2 \tag{3-4}$$

where: $L_1 = \beta_0 \tag{3-5a}$

$L_2 = 2\beta_1 - \beta_0 \tag{3-5b}$

$L_3 = 6\beta_2 - 6\beta_1 + \beta_0 \tag{3-5c}$

$L_4 = 20\beta_3 - 30\beta_2 + 12\beta_1 - \beta_0 \tag{3-5d}$

and, where the data ($x_{1:n}$) are first ranked in ascending order from 1 to n and:

$$\beta_0 = n^{-1} \sum_{j=1}^n x_j \tag{3-5e}$$

$$\beta_1 = n^{-1} \sum_{j=2}^n x_j [(j-1)/(n-1)] \tag{3-5f}$$

$$\beta_2 = n^{-1} \sum_{j=3}^n x_j [(j-1)(j-2)]/[(n-1)(n-2)] \tag{3-5g}$$

$$\beta_3 = n^{-1} \sum_{j=4}^n x_j [(j-1)(j-2)(j-3)]/[(n-1)(n-2)(n-3)] \tag{3-5h}$$

3-1.1 Graphical Depictions of L-Cv and L-Skewness

L-Cv is a dimensionless measure of variability. For a distribution or sample data that only has positive values, L-Cv is in the range $0 \leq L-Cv < 1$. Negative values of L-Cv are only possible if the at-site mean has a negative value. Table 3-1 lists some guidelines for describing the relative magnitude of L-Cv for a dataset or distribution.

L-Skewness is a dimensionless measure of asymmetry, which may take on positive or negative values. For a distribution or sample data, L-Skewness is in the range $0 \leq |L-Skewness| < 1$. Table 3-2 lists some guidelines for characterizing the relative magnitude of L-Skewness for a dataset or distribution.

Table 3-1 – General Descriptions of Relative Magnitude of L-Cv

ASSESSMENT OF THE RELATIVE MAGNITUDE OF L-Cv	
$.000 < L-Cv \leq .025$	minimal variability – often found in controlled manufacturing processes
$.025 < L-Cv \leq .075$	minor variability
$.075 < L-Cv \leq .150$	moderate variability
$.150 < L-Cv \leq .400$	large variability – often accompanied by large skewness
$.400 < L-Cv $	very large variability – often accompanied by very large skewness

Table 3-2 – General Descriptions of Relative Magnitude of L-Skewness

ASSESSMENT OF THE RELATIVE MAGNITUDE OF L-SKEWNESS	
L-skewness = 0.0	symmetrical distribution
$.000 < L-skewness \leq .050$	minor skewness
$.050 < L-skewness \leq .150$	moderate skewness
$.150 < L-skewness \leq .300$	large skewness
$.300 < L-skewness $	very large skewness, suggestive of “volatile” or outlier prone distributions

Figures 3-1a,b,c,d,e,f portray a number of frequency histograms to provide a visual depiction of how the magnitudes of L-Cv and L-Skewness relate to the shapes of sample datasets.

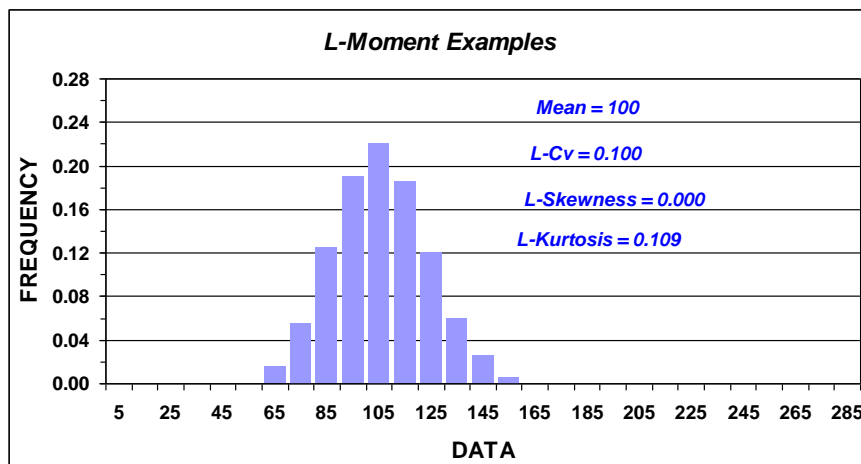


Figure 3-1a – Example Frequency Histogram for Small L-Cv for a Symmetrical Dataset




Figure 3-1b – Example Frequency Histogram for Moderate L-Cv for a Symmetrical Dataset




Figure 3-1c – Example Frequency Histogram for a Dataset with Moderate Negative L-Skewness




Figure 3-1d – Example Frequency Histogram for a Dataset with Moderate Positive L-Skewness




Figure 3-1e – Example Frequency Histogram for a Dataset with Large L-Skewness

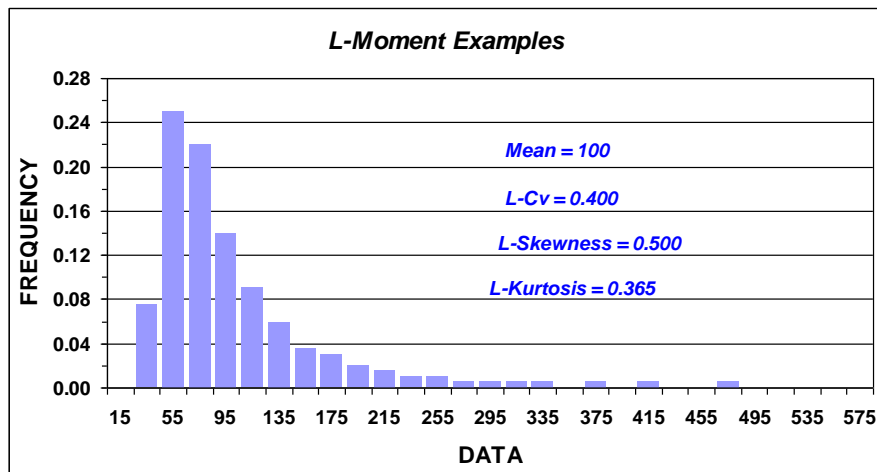


Figure 3-1f – Example Frequency Histogram for a Dataset with Very Large L-Cv and Very Large L-Skewness

3-1.2 Effect of Changes in L-Cv and L-Skewness on Regional Growth Curve

As described in Chapter 2, a homogeneous region satisfies the condition that all sites within the region can be described by one probability distribution having common distribution parameters after the site data are rescaled by their at-site mean. This rescaled non-dimensional probability distribution is termed a regional growth curve.

The conditions for a homogeneous region also equates to all sites having common values of the L-moment ratio values for L-Cv and L-Skewness. It is therefore useful to examine how changes in L-Cv and L-Skewness affect the behavior of the regional growth curve. An example was created for the Generalized Extreme Value (GEV) distribution to show the general behavior of how changes in L-Cv and L-Skewness affect the shape of the regional growth curve. Similar behavior would be seen for other 3-parameter probability distributions. Review of Figure 3-2a shows that if L-Cv is varied for a fixed value of L-Skewness, that changes in L-Cv affect the slope of the regional growth curve.

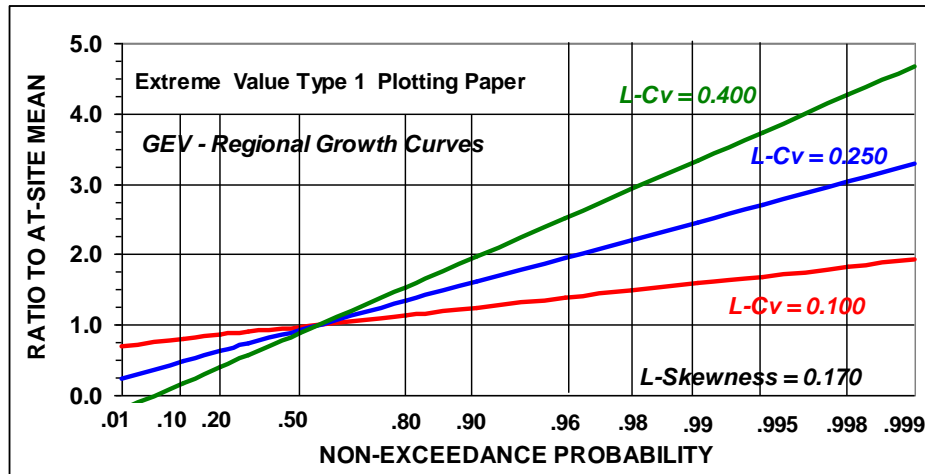


Figure 3-2a – Effect of Changes in L-Cv on Regional Growth Curve for GEV Distribution

Review of Figure 3-2b shows that changes in L-Skewness affect the shape (concave-convex) of the regional growth curve. For positive L-Skewness, increases in the magnitude of L-Skewness primarily affect the upper tail of the regional growth curve. Probability distributions with large positive L-Skewness are often said to be “volatile” having large quantile values for the upper tail. For the case of negative L-Skewness, increases in the magnitude of L-Skewness primarily affect the lower tail of the regional growth curve. This behavior should be kept in mind when considering how quantile estimates are affected by changes in the magnitude of L-Cv and L-Skewness.

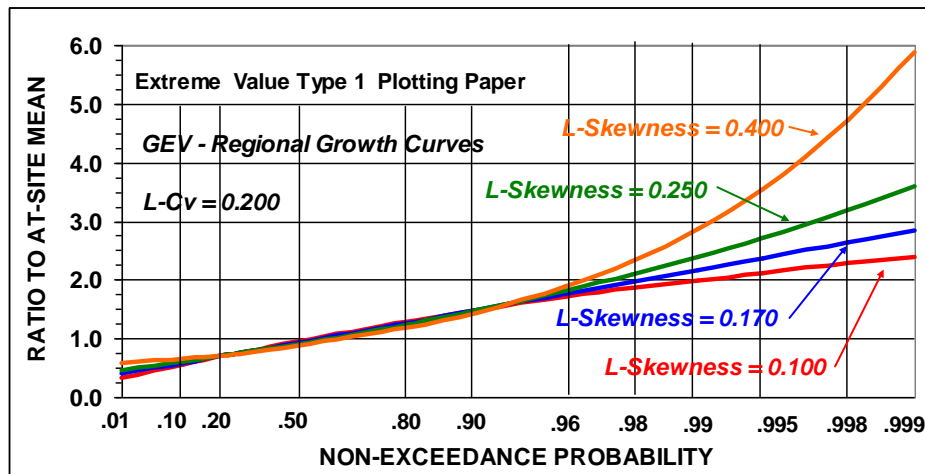


Figure 3-2b – Effect of Changes in L-Skewness to Regional Growth Curve for GEV Distribution

3-2 IDENTIFICATION OF DISCORDANT SITES WITHIN A GROUP OF SITES

In the process of grouping sites for proposed homogeneous regions, it is standard practice to compute a discordancy measure (D_i) for each site (Hosking and Wallis¹⁴, pp45-53). The discordancy measure is used to assist in identifying those sites whose L-moment ratios are discordant (markedly different) relative to L-moment ratios for the collection of sites. Specifically, sites with a discordancy measure greater than 3 are considered discordant relative to the collective behavior for the proposed grouping of sites.

There are two applications for use of the discordancy measure. The primary application is in the data screening process where the discordancy measure is used to identify suspicious datasets (sites) where data quality problems may be responsible for the discordant behavior. The second application is in conducting a regional analysis. In this application, if the proposed region is found to be heterogeneous as indicated by a large value of the heterogeneity measure (H1), then the physical characteristics of discordant sites may be helpful in understanding the cause of heterogeneity and assist in determining the course of action needed to produce a homogeneous region.

3-2.1 Discordant Sites Found in Data Screening Process

Those datasets that are flagged as discordant in the data screening process should be reviewed by examining the probability-plot and time-series graphics (data-screening tab) to determine if there are one or more data values whose magnitude differs markedly from the general behavior of the dataset. If discordant data values are found, the record for that site should be examined and a determination made if the data values are the result of some error in the measurement, recording or data entry process or if the values are valid. This is usually done by comparison of values with nearby sites or by using other independent information.

Identification of a discordant site or sites does not necessarily mean the site does not belong with the proposed grouping of sites. Rather, it is an indicator that additional research is needed to determine the cause of the discordancy before a decision can be made whether to keep the discordant site with the proposed group of sites; move the site to another grouping of sites; or remove the site from all analyses.

There are many possible reasons for a discordant site. Table 3-1 contains a partial list of possible reasons and recommended actions. In particular, it is important to note that low or high outliers should not be removed from the datasets. Low and high outliers at some sites are an expected outcome in large samples from multiple sites. The “apparent” outliers are important indicators about the natural variability of the phenomenon and the frequency of occurrence of low or high values. The real issue is the validity of the low or high values. If the data are valid - keep them. If the data are found to be erroneous, remove them.

Table 3-1 – Guidelines for Assessing and Handling of Discordant Sites During Data Screening

CAUSE OF DISCORDANCY	POSSIBLE ACTIONS
one or more data are suspicious, possibly due to measurement or recording errors	attempt to corroborate suspicious data from other sites or other sources of information, remove invalid data from dataset and re-compute discordancy measure
site dataset is generally of poor quality, numerous reasons possible	if data are found to be of poor quality, strong consideration should be given to removing site from all analyses
site dataset is small and an unrepresentative sample has occurred that can be attributed to sampling variability associated with small dataset	keep this site with proposed group of sites unless there are other considerations that cause doubt about inclusion of this site
site dataset includes one or more unusually large (small) values that can be verified as valid values, discordancy can be attributed to sampling variability	keep this site with proposed group of sites unless there are other considerations that cause doubt about inclusion of this site
no obvious cause of discordant measure	keep this site with proposed group of sites and delay decision until heterogeneity measure (H1) is computed during regional analyses. Conduct regional analyses with and without suspect site and compare magnitude of heterogeneity measure H1. Reassign site to another region if heterogeneity measure and other considerations indicate the site was incorrectly assigned initially.

3-2.2 Discordant Sites Found in Conducting Regional Analyses

In conducting the regional frequency analysis, the heterogeneity measure (H1) is the primary indicator for accepting or rejecting a proposed region (grouping of sites). The discordancy measures for the various sites provide a secondary indicator to consider whether a discordant site should be moved to another region. Table 3-2 provides some rough guidelines to assist in decision-making about what action, if any, to take about a discordant site(s).

Table 3-2 – Guidelines for Handling of Discordant Sites Found During Regional Analyses

HETEROGENEITY MEASURE	POSSIBLE ACTIONS
$H1 \leq 2$	No compelling reason to reassign discordant site unless large magnitude of discordancy at a particular site is interpreted to mean that the site is improperly assigned
$2 < H1 \leq 3$	Judgment required whether to accept as homogeneous region or to reassign discordant site(s) to another region to improve homogeneity of current grouping of sites
$3 < H1$	Proposed region is likely heterogeneous. Examine physical characteristics of discordant sites to assist in understanding of cause of heterogeneity relative to other sites. Consider alternative region formulation by reassigning sites to other regions.

3-3 HETEROGENEITY MEASURE

Heterogeneity of a proposed region is computed based on the magnitude of the site-to-site variability in the L-moment ratios relative to the level of variability expected in a homogeneous region (Hosking and Wallis¹⁴, pp61-63). The heterogeneity measures H1, H2 and H3 are for the L-moment ratios L-Cv, L-Skewness and L-Kurtosis, respectively. In practice, the H1 measure for the observed variability in L-Cv has been found to be a very useful measure. Conversely, the high level of natural variability in sample values of L-Skewness and L-Kurtosis result in the H2 and H3 measures having low discriminatory power. Therefore, the heterogeneity measure H1 for the level of variability in at-site values of L-Cv becomes the de-facto measure for assessing the relative level of heterogeneity for the proposed region.

The heterogeneity measure H1 is computed as follows. A weighted-average regional value of L-Cv ($L-Cv^R$) is computed from the sample values of L-Cv for the proposed sites in the region, where the weights are based on record length. A weighted-average standard deviation is then computed for the at-site L-Cv values for the collection of sites (V_{LCv}). A four-parameter Kappa distribution is fitted using the weighted-average regional values of the L-moment ratios for the sites. 500 computer simulations are then conducted using the fitted 4-parameter Kappa distribution, where each simulation has the same number of sites and record lengths as that for the proposed region. The mean (μ_v) and standard deviation (σ_v) are computed of the 500 samples of the standard deviation of the at-site samples of L-Cv. H1 is then computed as:

$$H1 = \frac{(V_{LCv} - \mu_v)}{\sigma_v} \quad 3-6$$

An H1 value of zero indicates that the site-to-site variability in at-site L-Cv values for the region are the same as would be expected from a homogeneous region with the observed L-moment ratios as fitted by a four-parameter Kappa distribution. Positive values of H1 indicate the site-to-site variability of at-site L-Cv values is greater than expected for a homogeneous region, and larger values of H1 indicate possible or likely heterogeneity (see Table 3-3). Conversely, negative values of H1 indicate the site-to-site variability of at-site L-Cv values is less than expected and the proposed region would be accepted as homogeneous.

An H1 value of 1.0 was originally proposed in Hosking and Wallis¹⁴ (pp63) for determining if the proposed region was acceptably homogeneous. That criterion was based solely on statistical considerations of the sampling characteristics for L-Cv based upon Monte Carlo studies with samples from known distributions. Oftentimes, there is additional variability in L-Cv that arises from difficulties in the accurate measurement and recording of data. In addition, there may be a variety of data quality control issues associated with human intervention in collecting and managing the data. Experience to date indicates that an H1 value of 2.0 is a reasonable choice for distinguishing between likely homogeneous and likely heterogeneous regions (Wallis³⁹). Table 3-3 lists guidelines for acceptance of proposed homogeneous regions where external sources add to the variability in L-Cv beyond statistical considerations of sampling variability. Adjustment of the threshold values of H1 in Table 3-3 may be warranted for analyses of some phenomenon based on the magnitude of variability imparted to sample data from consideration of data accuracy and quality control considerations. Somewhat lower threshold values of H1 may be used for environmental data where there is reasonably high accuracy in data measurement and recording and where good data quality control is possible.

Table 3-3 – Guidelines for Acceptance/Rejection of Proposed Region

HETEROGENEITY MEASURE	DECISION ON HOMOGENEOUS REGION
$H1 \leq 2$	Proposed Region is Acceptably Homogeneous
$2 < H1 \leq 3$	Marginal Heterogeneity Reassignment of some sites may be beneficial
$3 < H1$	Proposed Region is Likely Heterogeneous Reassignment of some sites is needed

3-4 DISTRIBUTION GOODNESS-OF-FIT MEASURE

A goodness-of-fit measure (Hosking and Wallis¹⁴, pp78-85) is used for identifying the probability distribution(s) that most closely matches the weighted-average regional L-Skewness and L-Kurtosis values for the grouping of sites for the proposed region. The test can best be visualized using the L-Moment Ratio Diagram shown in Figure 3-3. The test can be viewed as plotting the regional L-Skewness and L-Kurtosis values obtained from the sites for the proposed region and identifying the probability distribution(s) whose L-Kurtosis value most closely matches the regional L-Kurtosis value. The computation procedures for the goodness-of-fit measure Z_{dist} are described in detail on pp79-81 of Hosking and Wallis¹⁴.

3-4.1 Probability Distributions Available in L-RAP

Seven probability distributions are available in L-RAP for fitting of regional data. These seven distributions cover a wide portion of the L-Moment Ratio Diagram that has been found to be useful for describing environmental data. The seven distributions include:

- Generalized Logistic (GLO), special case of Kappa distribution with shape parameter $h = -1$
- Generalized Extreme Value (GEV), special case of Kappa distribution with shape parameter $h = 0$
- Generalized Normal (GNO)
- Gaucho, special case of Kappa distribution with shape parameter $h = +0.5$
- Generalized Pareto (GPA), special case of Kappa distribution with shape parameter $h = +1$
- Pearson 3 (P3)
- Kappa distribution (KAP)




Figure 3-3 – Example L-Moment Ratio Diagram for Selecting Best-Fit Probability Distribution

Figure 3-4 depicts an example of regional L-Skewness and L-Kurtosis pairings for 33 regions in Eastern Oregon for 24-hour precipitation annual maxima (Schaefer et al³⁶). Visually, it is apparent that the 24-hour precipitation annual maxima are well-described by Generalized Extreme Value (GEV) distribution. The numerical goodness-of-fit measures for the 33 regions confirmed the suitability of the GEV distribution.

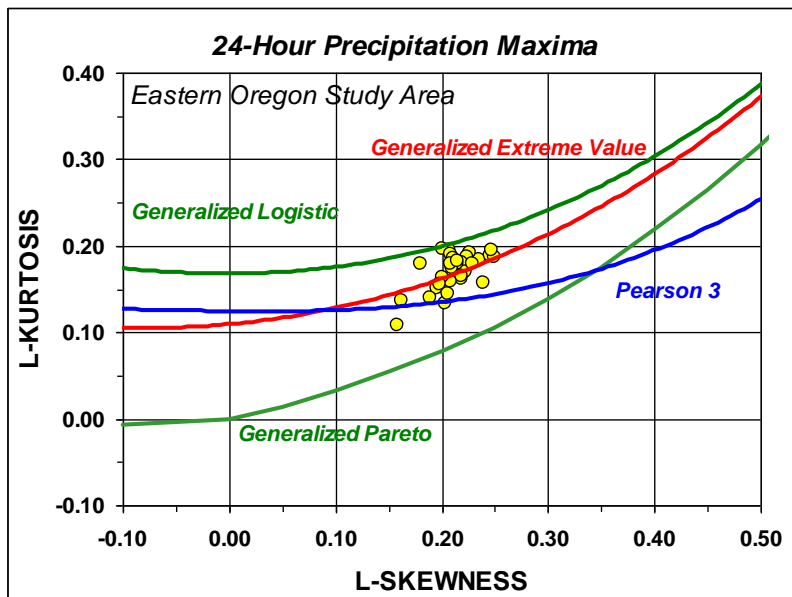


Figure 3-4 – Example L-Moment Ratio Diagram for Selecting Best-Fit Probability Distribution

3-4.2 Mixed Distributions

Some environmental datasets are comprised of a mixture of zero values and non-zero values. Other environmental data sets, such as water quality data, may be comprised of the number of samples below some detection threshold and non-zero values above the threshold. Both of these situations are best handled in a regional analysis framework by use of a mixed distribution (Hosking and Wallis¹⁴, pp76-77).

L-RAP includes procedures that counts the number of zero values and computes the proportion of zero values. Procedures are also included in L-RAP that provides quantile estimates for the mixed distribution where the probability distribution for the non-zero values is selected from the seven distributions listed above. The mixed distribution is described in Chapter 4 and all of the procedures for applying mixed distributions are described in Chapter 6.