WATER RESOURCES RESEARCH, VOL. 34, NO. 12, PAGES 3445–3459, DECEMBER 1998

The regional persistence and variability of annual streamflow in the United States

Richard M. Vogel, Yushiou Tsai, and James F. Limbrunner

Department of Civil and Environmental Engineering, Tufts University, Medford, Massachusetts

Abstract. Inference from individual streamflow records can be extremely misleading, even for large samples. One is often tempted to trust information available from a streamflow record rather than to exploit regional average statistics of those records. This study documents that regional average streamflow statistics usually contain much more information about the variability and persistence of streamflow at a particular site than does the individual streamflow record for that site. Experiments are performed using time series of annual streamflow at 1544 gauging stations across the continental United States. We document that 18 broad water resource regions of the United States are homogeneous in terms of the year-to-year persistence of streamflow, whereas much smaller regions are required to obtain homogeneity in terms of the variability of streamflow. Classical homogeneity measures ignore the serial correlation of streamflow. Instead, homogeneity is quantified using the sampling properties of at-site estimates of the coefficient of variation C_v and lag-one correlation ρ_1 of annual streamflows. Additional experiments using the Hurst coefficient reveal that the long-term persistence structure of historical annual streamflow series is indistinguishable from the long-term persistence structure of either an AR(1) or ARMA(1,1) process. If historical flow series are generated from either an AR(1) or ARMA(1,1) process, then even given 1544 observed time series, we are unable to distinguish between those two processes.

1. Introduction

Recent research in flood frequency analysis has documented the remarkable value of using regional hydrologic information, in addition to at-site data, for estimating the magnitude and frequency of flood events. Research and practice in this area have evolved to the point where practitioners now routinely employ regional methods such as hydrologic regression [*Jennings et al*., 1994]. *Bobée and Rasmussen* [1995] and *Hosking and Wallis* [1997] review recent advances in regional hydrologic methods applied to flood frequency analysis. Still, hydrologists are reluctant to use regional methods for other problems such as calibrating rainfall-runoff models, computing water balances, performing regional assessments, or investigating climate change. Before regional methods pervade other areas of hydrologic engineering we must first improve our understanding of the spatial and temporal structure of hydrologic data.

The primary goal of this study is to explore the regional and stochastic structure of annual flow records with an emphasis on characterizing the variability and persistence of those flow records for use in regional stochastic streamflow models. Other studies have examined the variability and persistence of monthly [*Bartlein*, 1982; *Guetter and Georgakakos*, 1993; *Lins*, 1997] and annual flow records [*Lins*, 1985a, b] in the United States using the multivariate statistical method known as principal components analysis (PCA) or eigenvector analysis. *Lins* [1985a, b] used PCA to evaluate the spatial and temporal structure of 48 sequences of annual flows in the United States. In those studies, PCA was used to identify the principal components of streamflow that explain its spatial and temporal structure. *Lins* [1985a, b] found that five principal components were able to explain most of the gross variations in annual streamflow across the United States and that most principal component functions were well described by an AR(1) process. In a more recent study, *Lins* [1997] used orthogonally rotated PCA and 559 monthly streamflow records in the United States to investigate the seasonal and persistence structure of streamflow. *Lins* [1997] documents that orthogonally rotated PCA can provide a useful approach to classifying regions of streamflow homogeneity in terms of both its temporal and spatial structure.

This study differs significantly from the studies cited above because we attempt to explore the stochastic structure of the annual flow series directly, rather than focus on the structure of their principal components or eigenvectors. Regional examinations of annual flow series allow us to select a suitable regional stochastic model and simultaneously to obtain regional estimates of statistics that characterize the variability and persistence of the flow series. The goal of this study is to provide an initial basis for the development of regional stochastic time series models of annual streamflow for the conterminous United States.

There is significant literature which addresses methods for assessing the homogeneity of streamflow records, with particular emphasis on estimation of flood quantiles [*Hosking and Wallis*, 1997]. Since one of the objectives of this study is to examine the stochastic structure of flow records, existing methods for evaluating homogeneity which assume independence among individual flows [see, for example, *Fill and Stedinger*, 1995; *Hosking and Wallis*, 1997] could not be employed here. Our approach is to begin with an experiment that evaluates the regional and stochastic structure of streamflow persistence using a large database of historical streamflows for the United

Paper number 98WR02523.
0043-1397/98/98WR-02523$09.00

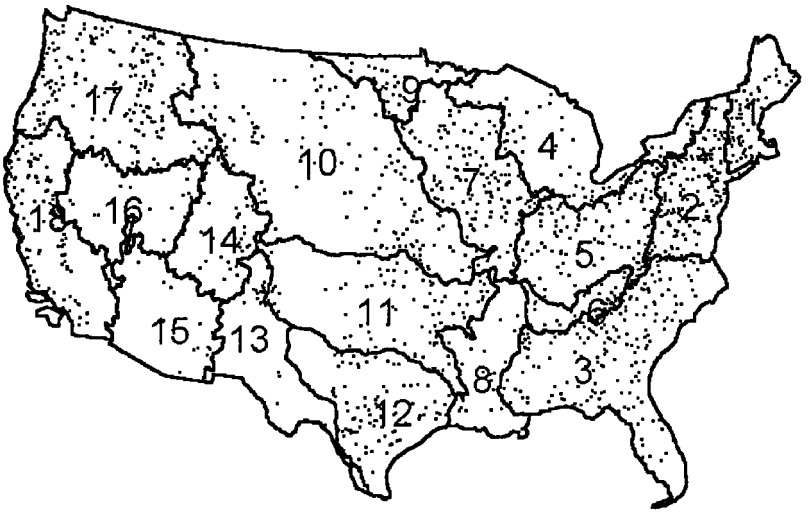

Figure 1. Locations of streamflow gauges [from *Slack et al*., 1993] within each of the 18 major U.S. water resource regions.

States. These experiments allow us to choose a suitable regional stochastic structure of flow series. Subsequent experiments exploit that regional stochastic structure to examine the regional variability (homogeneity) of streamflow.

1.1. Database of Streamflow

The data set consists of records of average annual streamflow taken from a climate-sensitive database of 1569 stream gage records available on CD-ROM and the World Wide Web (http://wwwrvares.er.usgs.gov/hcdn_cdrom/1st_page.html) from the U.S. Geological Survey (USGS) [*Slack et al*., 1993]. Dropping the records from Alaska, Hawaii, and Puerto Rico, as well as all those sites with observed annual flows lower than 0.283 L/s, left us with 1544 sites. Annual flows are computed using water years (October–September). The locations of streamflow gauging stations within each of 18 major water resource regions of the United States are illustrated in Figure 1. In 1970 the U.S. Water Resources Council defined 21 major geographic areas, or regions, for the purpose of assessing the state of water resources across the nation [*U.S. Water Resources Council*, 1975]. Dropping Alaska, Hawaii, and Puerto Rico leaves the 18 regions depicted in Figure 1.

The streamflow data set, termed the Hydroclimatologic Data Network (HCDN), is made up of USGS streamflow records which are relatively free of anthropogenic influences. Since details of the data set are discussed elsewhere, we do not repeat those discussions here [see *Lins*, 1997; *Vogel and Wilson*, 1996; *Slack et al*., 1993].

The experiments that follow involve comparisons among estimates of coefficient of variation C_v, lag-one correlations ρ_1, and the Hurst coefficient H taken from these flow records. For estimation of H we employ the longest continuous record at 1544 of the HCDN sites. However, for the statistics C_v and ρ_1, we employ a fixed streamflow record length of $n = 40$ years to ensure that differences in sampling variability observed for each flow statistic do not result from differences in record lengths. Our goal was to include as many flow records m, each with as long a record length n, as possible. Maximization of the product mn for this database resulted in 885 sites out of the original 1544 sites with record lengths of 40 years or more. The first 40 years of record were used for each site which had at least 40 years of record. The number of sites in each region with $n \geq 40$ is listed in Table 1.

1.2. The Assumed Stochastic Structure of Annual Streamflow

Markovic [1965], *Vogel and Wilson* [1996], and others summarize hypothesis tests and goodness-of-fit evaluations of various probability distributions to sequences of annual streamflow in the continental United States. On the basis of those evaluations, *Markovic* [1965] and *Vogel and Wilson* [1996] conclude that sequences of annual streamflow are well approximated by either a Gamma or two-parameter lognormal distribution. As anticipated, various three-parameter distributions such as the lognormal, log Pearson, or Pearson distributions are also suitable yet probably unnecessary. In this study we assume that annual streamflow sequences x_t, $t = 1, \ldots, n$ arise from an autoregressive lognormal process termed an AR(1)-LN model defined by

$$y_t = \mu_y + \rho_{y,1}(y_{t-1} - \mu_y) + \varepsilon_t \sigma_y \sqrt{1 - \rho_{y,1}^2} \tag{1}$$

where $y_t = \ln(x_t)$ and the moments of x are related to the moments of y via

$$\mu_y = \ln\left(\frac{\mu}{1 + C_v^2}\right) \tag{2}$$

$$\sigma_y = \sqrt{\ln(1 + C_v^2)} \tag{3}$$

$$\rho_{y,1} = \frac{\ln(\rho_1(e^{\sigma_y^2} - 1) + 1)}{\sigma_y^2} \tag{4}$$

$$C_v = \frac{\sigma}{\mu} \tag{5}$$

Here μ, σ, and ρ_1 are the mean, standard deviation, and lag-one serial correlation, respectively, of the flows and μ_y, σ_y, and $\rho_{y,1}$ are the mean, standard deviation, and lag-one serial correlation, respectively, of the natural logarithms of the flows. Experiments are performed later in this study to evaluate the goodness of fit of this model to annual flows in the United States. *Lins* [1985a, b] assumed a lognormal model of annual

Table 1. Summary of Regional Statistics of the Lag-One Serial Correlation of Annual Streamflow in the Continental United States

Region	Number of Streamflow Records (m)	Number of Records $n \geq 40$	Average Record Length ($\bar{n}$)	Regional Mean of Logs ($\bar{\rho}_{y1}$)	Regional Mean of Flows (r_1^*)	Regional Mean of Flows (r_2^*)
1	71	43	46.0	0.206	0.197	0.068
2	167	114	47.2	0.247	0.264	0.116
3	193	104	41.3	0.278	0.269	0.090
4	57	35	45.8	0.383	0.361	0.144
5	108	87	51.4	0.251	0.283	0.040
6	44	27	48.5	0.270	0.271	0.051
7	126	101	50.7	0.349	0.350	0.082
8	23	14	42.8	0.181	0.193	0.028
9	32	14	38.5	0.225	0.149	0.055
10	140	70	39.9	0.259	0.218	0.112
11	87	42	39.5	0.237	0.219	0.048
12	83	37	38.5	0.222	0.110	0.081
13	22	8	37.4	0.119	0.114	0.039
14	44	17	40.9	0.134	0.146	0.045
15	16	9	41.4	0.099	0.070	0.018
16	32	14	41.8	0.250	0.261	0.131
17	189	100	41.6	0.084	0.051	0.187
18	110	49	41.5	0.098	0.032	0.105
All regions	1544	885	43.0	0.22	0.198	0.080

streamflow in his studies of the interannual streamflow variability of streamflow in the United States.

2. Regional Short-Term Persistence of Streamflow

Analogous to the previous section, we now explore the theoretical and empirical structure of some simple measures of persistence. The most common index of short-term persistence is the lag-k correlation coefficient ρ_k, which is defined as

$$\rho_k = \frac{\text{Cov}[x_i, x_{i+k}]}{\text{Var}[x]} \tag{6}$$

It is a measure of short-term persistence because beyond a few lags k, the coefficient usually decays to noise. The most common index of short-term persistence, ρ_1, only quantifies the correlation between successive values of a random variable, as opposed to long-term measures of persistence such as the Hurst coefficient, discussed in a later section. One common estimator of the lag-k serial correlation coefficient ρ_k is

$$r_k = \frac{\sum_{i=1}^{n-k} x_i x_{i+k} - \frac{1}{n-k}\sum_{i=1}^{n-k} x_i \sum_{i=1}^{n-k} x_{i+k}}{\left[\sum_{i=1}^{n-k} x_i^2 - \frac{1}{n-k}\left(\sum_{i=1}^{n-k} x_i\right)^2\right]^{1/2}\left[\sum_{i=1}^{n-k} x_{i+k}^2 - \frac{1}{n-k}\left(\sum_{i=1}^{n-k} x_{i+k}\right)^2\right]^{1/2}} \tag{7}$$

In this section we concentrate on the serial correlation of the logarithms of streamflow $\rho_{y,k}$ because the sampling properties of an AR(1) normal process are available. The at-site estimator of $\rho_{y,k}$, termed $r_{y,k}$, is obtained from (7) by replacing x with y. To evaluate the properties of at-site estimates of $\rho_{y,k}$, we use the biased estimator in (7) because *Tasker* [1983] approximated its probability density function (pdf), assuming an AR(1) normal process, so that

$$f(r_{y,1}|\rho) = \frac{(1 - r_{y,1}^2)^{(1/2)(T-1)}(1 + c^2 - 2cr_{y,1})^{-(1/2)T}}{\beta[\frac{1}{2}, \frac{1}{2}(T+1)]} \tag{8}$$

where $T = n - 2$ and $c = \rho_{y,1} - [(1 + \rho_{y,1})/(n - 3)]$. We could have also used the exact pdf given by *Leipnik* [1947]; however, that result is for another biased estimator of $\rho_{y,1}$ whose small-sample properties have not been investigated. For large n, Leipnik's pdf is equivalent to Tasker's pdf; however, for $n = 40$, the pdfs are different.

The estimator of ρ_1 in (7) is used with (8) to derive the frequency distribution of estimates of ρ_1 in each region. To obtain unbiased estimates of ρ_1, we employ the estimator of ρ_1 given by equation (3) of *Wallis and O'Connell* [1972], which, like (7), is downward biased and, in particular for $k = 1$, has approximate expectation

$$E[r_1] = \rho_1 - \frac{1 + 4\rho_1}{n} \tag{9}$$

so that an unbiased estimator r_1^* can be derived from (9), resulting in

$$r_1^* = \frac{nr_1 + 1}{n - 4} \tag{10}$$

Stedinger [1981] shows that r_1^* is a nearly unbiased estimator of ρ_1 even for an AR(1)-LN model.

Figure 2 illustrates box plots of the unbiased estimator r_1^*, of annual streamflow, for each of the 18 water resource regions. The lag-one correlations are generally higher in the more temperate regions where groundwater outflow causes greater year-to-year persistence in streamflow. In arid and semiarid regions the lag-one correlations are nearly zero, indication that groundwater outflow plays much less of a role in producing surface water runoff.

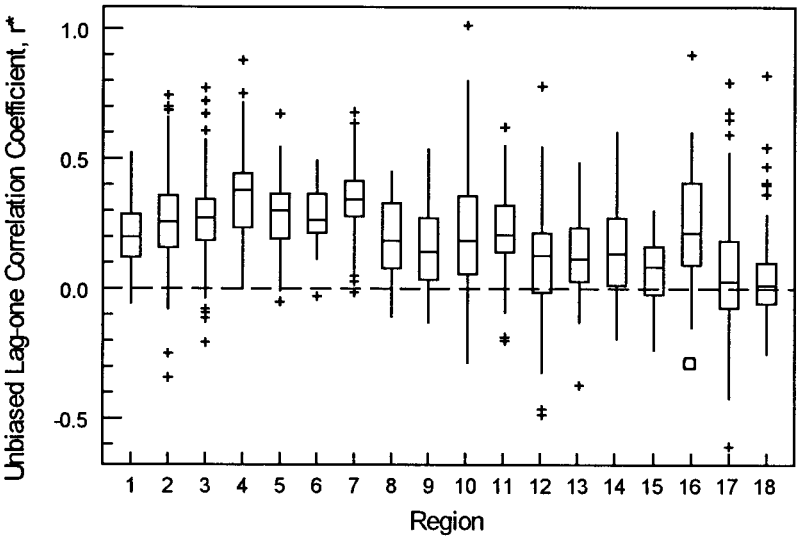

Figure 2. Box plots of the unbiased lag-one serial correlation coefficient r_1 for the 18 water resource regions.

An examination of Figure 2 led us to conjecture that each of the 18 water resource regions are homogeneous in terms of ρ_1. The goal of this section is to explore whether differences in estimates of ρ_1 within a region are a result of sampling variability or due to real differences in the physical processes that lead to streamflow persistence.

In the following experiments the lag-one correlations are based on the logarithms of streamflow, because the pdf's introduced by *Leipnik* [1947] and *Tasker* [1983] correspond to AR(1) normal inflows. Figure 3 compares histograms of estimates of $\rho_{y,1}$ for each region with the theoretical probability density function $f(r_{y,1}|\rho_{y,1})$ given in (8). In each region the true value of $\rho_{y,1}$ is assumed equal to an unbiased estimate of its regional mean value computed using the record length weighted estimator

$$\bar{r}_{y,1} = \frac{\sum_{i=1}^{m} n_i r^*_{y,1}(i)}{\sum_{i=1}^{m} n_i} \tag{11}$$

where n_i and $r^*_{y,1}(i)$ are the record length and unbiased estimate, respectively, of $\rho_{y,1}$ at site i, and m is the number of sites in the region. The regional mean value $\bar{r}_{y,1}$ in (11) is based on the complete set of 1544 streamflow records summarized earlier in Table 1.

The plots in Figure 3 document that it is entirely possible that the differences in estimates of $\rho_{y,1}$ within each region could be due to sampling alone. What appear at first to be rather large variations in short-term persistence within each region could be the result of sampling variability rather than differences in physical processes. This result emphasizes the importance of using regional information to understand physical processes such as short-term persistence, because each individual record is really only one realization from an ensemble of realizations from that region.

Figure 3 also reports results for the entire United States. Interestingly, if the regional mean value of $\rho_{y,1}$ is assumed to be 0.228 for the entire United States, then it is entirely possible that the observed variability in sample estimates of the lag-one correlation, across the entire nation, is due to sampling alone. This is not a surprising result because in 11 of the 18 regions of the United States, the regional value of $\rho_{y,1}$ reported in Tables 1 and 2 ranges from 0.181 to 0.278. It is only the north central regions 4 and 7 that exhibit significantly higher regional mean values of $\rho_{y,1}$, values and the southwestern regions 13, 14, 15, and 18 and northwestern region 17 that exhibit significantly lower regional mean values of $\rho_{y,1}$. We conclude that it is probable that there are roughly three different regions of roughly homogeneous short-term persistence. There is very low persistence in the southwestern and northwestern regions, high persistence in the north central region, and moderate persistence in all other regions of the United States.

3. The Regional Variability of Streamflow

The most common index of variability is the coefficient of variation C_v defined in (5). This section explores the regional homogeneity of C_v by comparing the theoretical behavior of sample estimates of C_v with the observed behavior of sample estimates of C_v for broad regions of the United States. Prior to such experiments it is necessary to understand the theoretical sampling properties of estimates of C_v. To our knowledge, there is no literature which examines the properties of estimates of C_v for correlated processes, and hence we begin by introducing and comparing several estimators.

It is common practice to estimate C_v using the ratio estimator

$$\hat{C}_{v1} = \frac{s}{\bar{x}} \tag{12}$$

where

$$\bar{x} = \frac{1}{n}\sum_{i=1}^{n} x_i \qquad s^2 = \frac{1}{n-1}\sum_{i=1}^{n} (x_i - \bar{x})^2$$

are standard sample estimators of the mean and variance. *Kirby* [1974] documents that $\hat{C}_{v1}$ is bounded above and *Vogel*

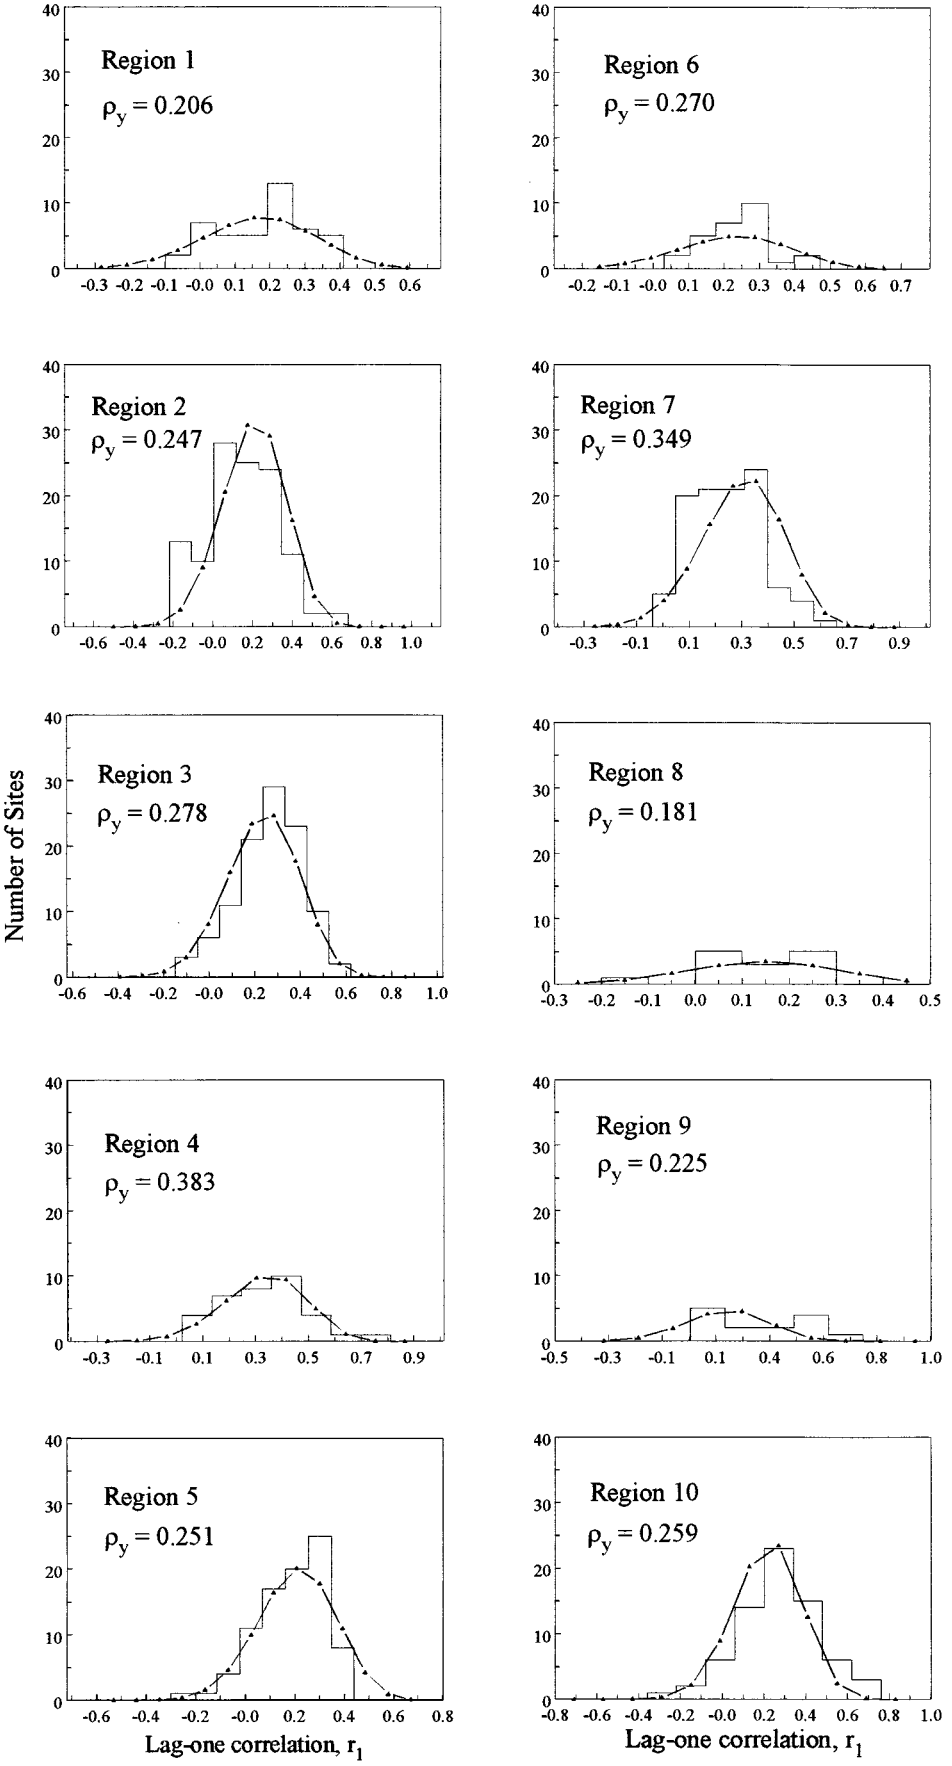

Figure 3. Comparison of histogram of r_1 with the theoretical probability density function of r_1.

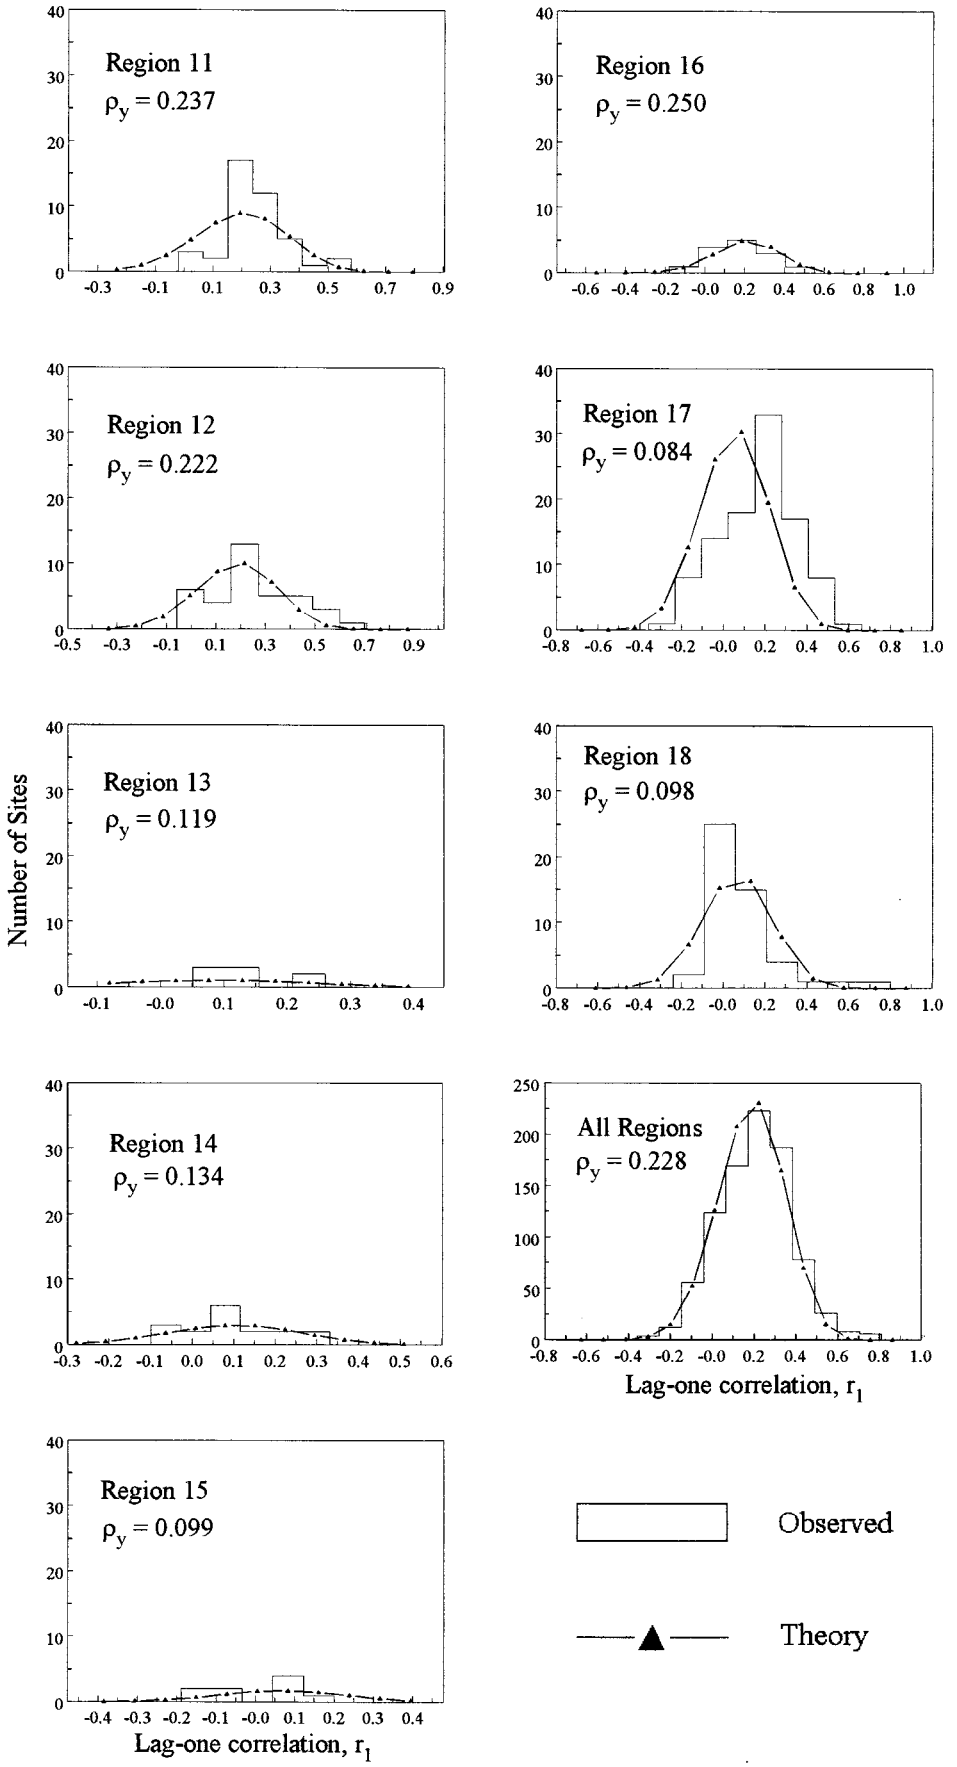

Figure 3. (continued)

Table 2. Summary of Regional Statistics of the Coefficient of Variation of Annual Streamflow in the Continental United States

Region	Number of Streamflow Records (m)	Average Record Length ($\bar{n}$)	$\bar{C}_v$	Median C_v	Theoretical ($C_v(\hat{C}_{v3})$) (From Equation (13))	Sample Estimate ($C_v(\hat{C}_{v3})$)	Regional Mean ($\bar{\rho}_{y1}$)
1	71	46.0	0.277	0.285	0.126	0.159	0.206
2	167	47.2	0.316	0.315	0.130	0.240	0.247
3	193	41.3	0.415	0.369	0.137	0.387	0.278
4	57	45.8	0.293	0.278	0.143	0.373	0.383
5	108	51.4	0.386	0.370	0.133	0.455	0.251
6	44	48.5	0.280	0.273	0.130	0.179	0.270
7	126	50.7	0.635	0.594	0.157	0.428	0.349
8	23	42.8	0.472	0.498	0.132	0.258	0.181
9	32	38.5	1.799	1.188	0.231	0.507	0.225
10	140	39.9	0.840	0.720	0.161	0.686	0.259
11	87	39.5	0.792	0.711	0.156	0.546	0.237
12	83	38.5	1.323	1.093	0.194	0.389	0.222
13	22	37.4	0.703	0.602	0.142	0.350	0.119
14	44	40.9	0.382	0.356	0.126	0.229	0.134
15	16	41.4	0.786	0.740	0.147	0.325	0.099
16	32	41.8	0.651	0.551	0.147	0.470	0.250
17	189	41.6	0.331	0.299	0.122	0.393	0.084
18	110	41.5	0.980	0.765	0.160	0.947	0.098
All regions	1544	43.0	0.59				0.22

and Fennessey [1993] document that this estimator is remarkably downward biased, with that bias increasing as the skewness increases and/or sample size decreases. An alternative estimator is based on (3), which can be rearranged to obtain

$$C_v = \sqrt{\exp(\sigma_y^2) - 1} \tag{13}$$

A natural sample estimator which results from (13) is

$$\hat{C}_{v2} = \sqrt{\exp(v_y^2) - 1} \tag{14}$$

where

$$v_y^2 = \frac{1}{n}\sum_{i=1}^{n}(y_i - \bar{y})^2 \qquad \bar{y} = \frac{1}{n}\sum_{i=1}^{n} y_i$$

with $y_t = \ln(x_t)$. However, even $\hat{C}_{v2}$ will be biased, because the estimator v_y^2 is known to be biased, especially for serially correlated variables. *Kendall* [1954] derived the expectation of v_y^2 for an AR(1) normal process

$$E[v_y^2] = \sigma_y^2\left[1 - \frac{1}{n} - \frac{2}{n}\sum_{k=1}^{n-1}\left[1 - \frac{k}{n}\right]\rho_{y,k}\right] \tag{15}$$

where $\rho_{y,k}$ is the lag-k autocorrelation of the y series. The downward bias associated with estimators v_y^2 and C_{v1} results, in part, from the correlation of the time series. Since this study attempts to understand the regional behavior of C_v, an unbiased estimator of C_v is needed, and hence we derive one here. The correlogram of an autoregressive process is $\rho_{y,k} = \rho_{y,1}^k$, where $\rho_{y,k}$ denotes the lag-k autocorrelation of y. Substitution of $\rho_{y,k} = \rho_{y,1}^k$ into (15), using the fact that (15) is an arithmetic power series, and dropping the ρ^n terms yields

$$E[v_y^2] \cong \sigma_y^2\left[1 + \frac{1}{n} + \frac{2[\rho_{y,1}(n+1) - n]}{n^2(1 - \rho_{y,1})^2}\right] \tag{16}$$

so an approximately unbiased estimate of σ_y^2 for an autoregressive process is

$$v_y^{*2} = \frac{v_y^2}{\left[1 + \frac{1}{n} + \frac{2[\rho_{y,1}(n+1) - n]}{n^2(1 - \rho_{y,1})^2}\right]} = \frac{v_y^2}{\theta} \tag{17}$$

Note that for an independent process (17) reduces to the expected result $v_y^{*2} = (n/(n-1))v_y^2$. A nearly unbiased estimator of C_v for an autoregressive lognormal process is then

$$\hat{C}_{v3} = \sqrt{\exp(v_y^{*2}) - 1} \tag{18}$$

Prior to application of these estimators on actual streamflow samples, we investigate their sampling properties using Monte Carlo experiments. Such experiments should assist us in separating out real physical differences in C_v from variations caused by sampling.

3.1. Sampling Properties of Estimators of the Coefficient of Variation

A Monte Carlo experiment was performed to compare the bias associated with the estimators $\hat{C}_{v1}$, $\hat{C}_{v2}$, and $\hat{C}_{v3}$. The experiment involved generating 50,000 samples of length $n = 40$ and 100,000 samples of length $n = 20$ from the AR(1)-LN model given in (1) for various values of C_v and ρ_1. Figure 4 illustrates the bias associated with $\hat{C}_{v1}$, $\hat{C}_{v2}$, and $\hat{C}_{v3}$, where bias $(\hat{C}_v) = C_v - E[\hat{C}_v]$. In general, both estimators $\hat{C}_{v1}$ and $\hat{C}_{v2}$ are downward biased, and that bias increases as n decreases, as C_v increases, and as ρ_1 increases. The estimator $\hat{C}_{v3}$ is generally unbiased, and for this reason it is used in the remainder of this study to obtain an unbiased estimate of regional mean values of C_v. Note that $\hat{C}_{v3}$ is only unbiased when the true value of $\rho_{y,1}$ is known, as was assumed in these experiments. As is shown in the previous section, unbiased regional estimates of $\rho_{y,1}$ provide a very good approximation to the true value of $\rho_{y,1}$, and hence, for the purposes of this study, $\hat{C}_{v3}$ turns out to be a very appropriate and useful estimator. However, for other purposes, it may be preferable to employ $\hat{C}_{v2}$, which exhibits larger bias yet lower variance than $\hat{C}_{v3}$.

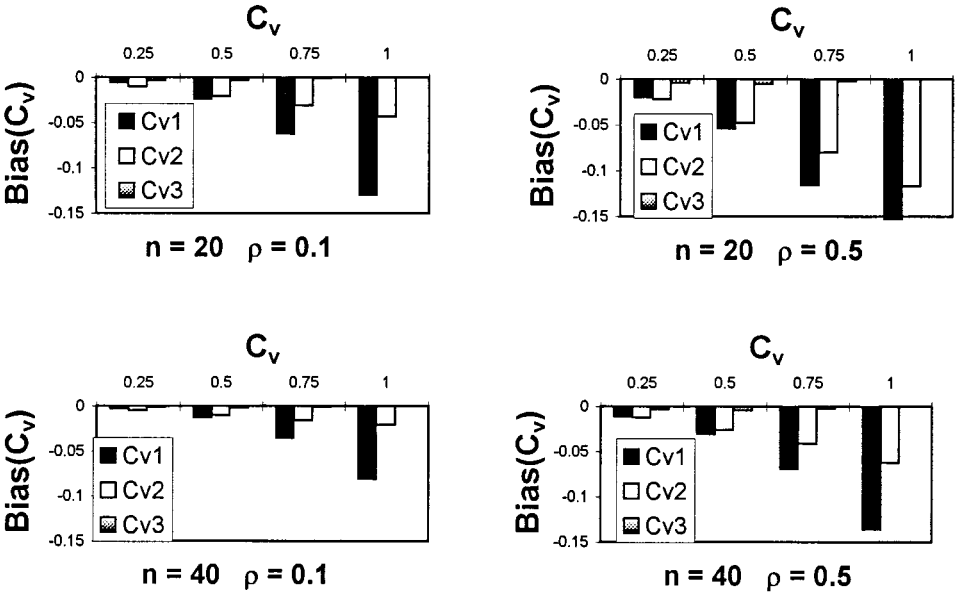

Figure 4. Bias associated with $\hat{C}_{v1}$, $\hat{C}_{v2}$, and $\hat{C}_{v3}$ for an AR(1)-LN process.

Another goal of this study is to explore the sampling distribution of $\hat{C}_{v3}$ so that we can evaluate whether estimates of C_v in a region are due to sampling variability or due to real differences in hydrology. One approach to examining the variability of estimates of C_v is to derive an analytical expression for the Var$[\hat{C}_{v3}]$. *Stuart and Ord* [1987, equation 10.19] derive a first-order approximation to Var$[\hat{C}_{v1}]$ for an independent normal process. *Fill and Stedinger* [1995] derive a first-order approximation to Var$[\hat{C}_{v1}]$ for any distribution. A first-order approximation to Var$[\hat{C}_{v3}]$ for an AR(1)-LN process is derived in the appendix, resulting in

$$\text{Var}[\hat{C}_{v3}] \cong \frac{[(1 + C_v^2)\ln(1 + C_v^2)]^2}{2n(\theta C_v)^2}\left[\frac{1 + \rho_{y1}^2}{1 - \rho_{y1}^2}\right] \qquad (19)$$

where

$$\theta = \left[1 + \frac{1}{n} + \frac{2[\rho_{y,1}(n + 1) - n]}{n^2(1 - \rho_{y,1})^2}\right].$$

(Figure A1 also documents that the approximation in (19) compares favorably with Var$[\hat{C}_{v3}]$ estimated using a Monte Carlo experiment.)

3.2. Experiment 1: Regional Homogeneity of Streamflow

An experiment is performed to determine whether or not the 18 major water resource regions in Figure 1 are homogeneous in terms of the year-to-year variability of streamflow. We address this issue by determining whether or not it is possible that C_v is fixed (except for sampling variability) in each of the 18 regions in Figure 1. Estimates of $\hat{C}_{v3}$ are obtained for each of the 1544 basins in Figure 1. Figure 5 summarizes box plots of sample estimates of $\hat{C}_{v3}$ for each region using solid lines. For comparison, we also illustrate in Figure 5 (using dashed lines) theoretical box plots based on the assumption that $\hat{C}_{v3}$ follows a lognormal distribution with Var$[\hat{C}_{v3}]$ given by (19) and the mean value of $\hat{C}_{v3}$ given by (20), below. To compute $\hat{C}_{v3}$ at each site, and Var$[\hat{C}_{v3}]$ for the region, the true value of $\rho_{y,1}$ is required. We use the unbiased estimates of the regional mean value of $\rho_{y,1}$ in (11) (denoted $\bar{\rho}_{y1}$), which are summarized in Tables 1 and 2. The box plots of the sample estimates of $\hat{C}_{v3}$ in Figure 5 are all based on estimates of $\hat{C}_{v3}$ derived from samples of equal length $n = 40$. For the theoretical box plots illustrated in Figure 5 and the regional mean values reported in Table 2, we use all streamflow samples in the region to estimate the regional mean value of $\hat{C}_{v3}$, denoted $\bar{C}_v$, computed using the record length weighted estimator

$$\bar{C}_v = \frac{\sum_{i=1}^{m} n_i \hat{C}_{v3}(i)}{\sum_{i=1}^{m} n_i} \qquad (20)$$

where n_i and $\hat{C}_{v3}(i)$ are the record length and unbiased estimate of C_v, respectively, at site i, and m is the number of sites in the region. The average regional values of C_v and $\rho_{y,1}$ reported in Table 2 are based on the complete set of 1544 basins. Also shown in Table 2 are the number of sites in each region m and the average record length $\bar{n}$ in each region for the sites used to compute the regional mean values of C_v and $\rho_{y,1}$ and the median value of C_v for each region.

Table 2 also compares the theoretical and observed coefficient of variation of the estimated values of $\hat{C}_{v3}$. Here the coefficient of variation of $\hat{C}_{v3}$ was computed as $C_v(\hat{C}_{v3}) = \sqrt{\text{Var}[\hat{C}_{v3}]}/\bar{C}_v$, with the theoretical value of Var$[\hat{C}_{v3}]$ computed using (19) and the sample estimate of $C_v[\hat{C}_{v3}]$ obtained using $\hat{C}_{v1}[\hat{C}_{v3}]$. A comparison of the values of $C_v(\hat{C}_{v3})$ in Table 2 reveals that the observed values of C_v are uniformly more variable than one would expect due to sampling alone. This same result may be viewed graphically in Figure 5, which illustrates that the theoretical box plots are generally less variable than are the box plots of the observed at-site sample estimates. In fact, Figure 5 illustrates that only regions 1, 6, and possibly 14 are nearly homogeneous in terms of C_v because the theoretical and observed box plots are nearly equivalent. Therefore we conclude, similar to *Lins* [1997], that the 18 water resource regions are not homogeneous in terms of the year-to-year variability in streamflow.

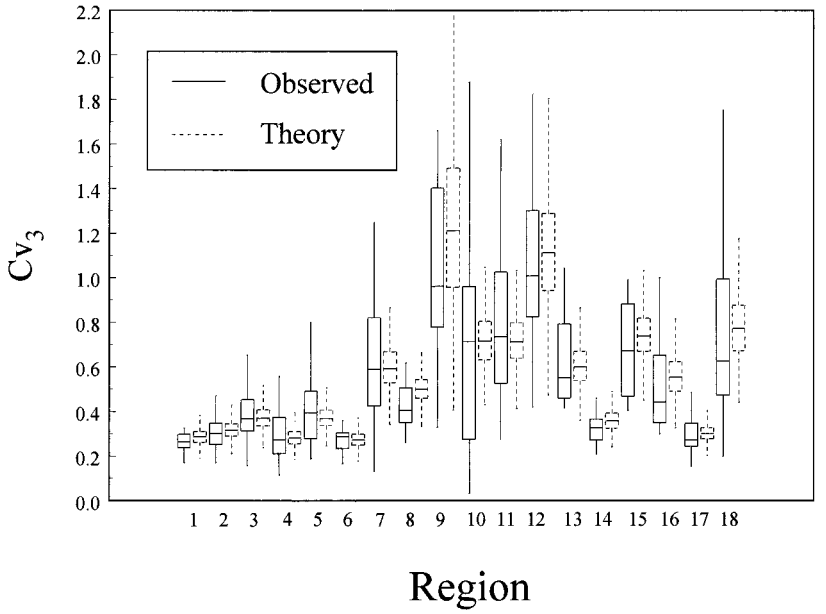

Figure 5. Comparison of observed and theoretical box plots of estimates of C_v for 18 water resource regions.

In Figure 6 we explore whether or not breaking the 18 regions into smaller subregions yields more homogeneity in C_v. Figure 6 compares theoretical and observed box plots of $\hat{C}_{v3}$ for selected subregions within regions 1, 4, 10, and 18. Here subregions are defined using the standard 222 subregions defined by the *U.S. Water Resources Council* [1975]. While, in general, Figure 6 documents that splitting regions into smaller subregions leads to more homogeneity in terms of C_v, the results are still inconclusive. When splitting regions into smaller subregions, there were often too few sites left in each region to obtain reliable box plots. For example, the observed box plots shown in Figure 6 are based on between 4 and 14 sites, with an average of only six sites per subregion. In very heterogeneous regions such as 4, 10, and 18, Figure 6 documents that splitting into subregions leads to much greater homogeneity of each subregion. Unfortunately, there are not enough data in each subregion to draw conclusions any more definitive than this.

4. Regional Long-Term Persistence of Streamflow

Ever since *Hurst* [1951] introduced a methodology for exploring the long-term persistence of geophysical records, nu-

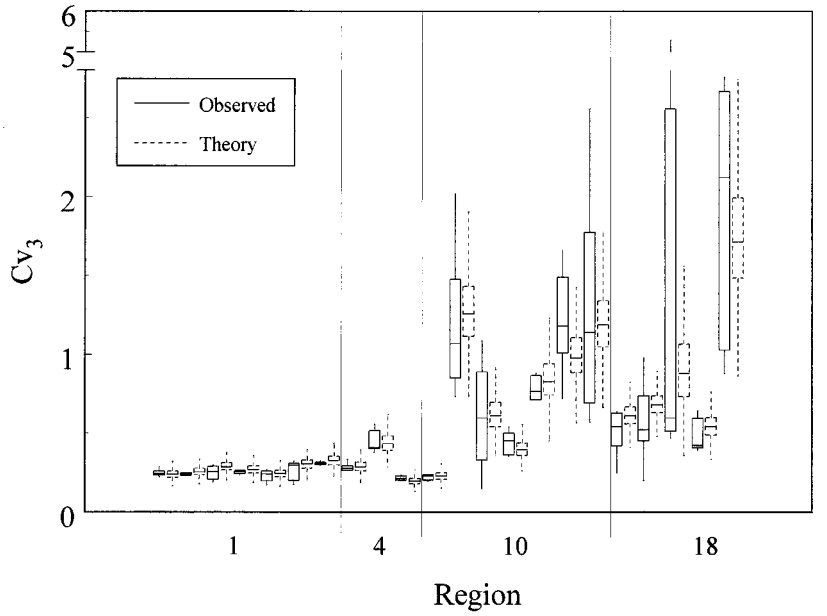

Figure 6. A comparison of observed and theoretical box plots of estimates of C_v for all subregions within regions 1, 4, 10, and 18.

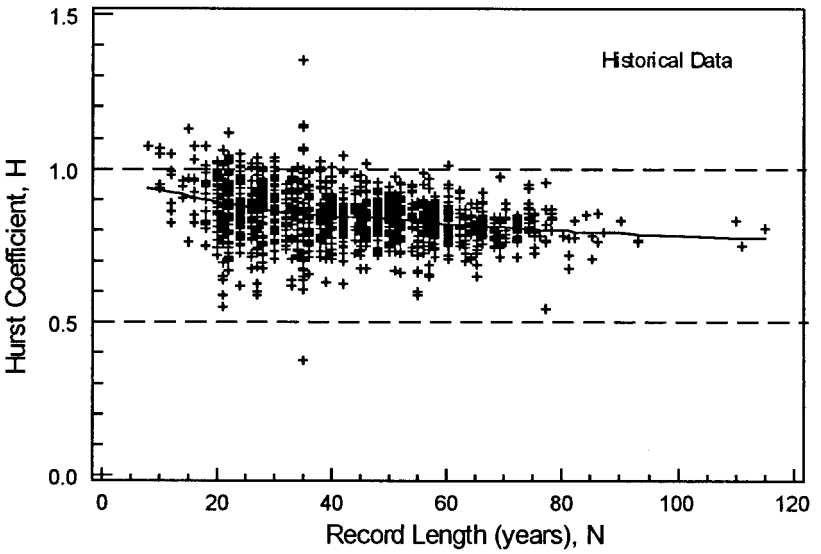

Figure 7. Estimates of Hurst coefficients of annual streamflow from 1544 historical records in the continental United States.

merous investigators have used his approach to explore the stochastic structure of streamflow, climate, and other geophysical records. The Hurst coefficient H is defined as follows: Given a streamflow series $\{x_1, x_2, \cdots, x_n\}$ with sample mean $\bar{x}$ and sample variance s_x^2, the adjusted partial sums are

$$S_k = \sum_{t=1}^{k} x_t - k\bar{x} \qquad k = 1, \cdots, n \tag{21}$$

and the range is

$$R_n = [\max(S_1, S_2, \ldots, S_n) - \min(S_1, S_2, \ldots, S_n)] \tag{22}$$

Hurst [1951] found that

$$E\left[\frac{R_n}{s_x}\right] \propto n^H \tag{23}$$

The exponent H, termed the Hurst coefficient, was found to vary between 0.69 and 0.80, whereas independent normal variables are known to have $H = 0.5$. Even a simple AR(1) model will exhibit $H = 0.5$ for very large n. These findings led many investigators to speculate and introduce more complex stochastic models, which purport to reproduce the Hurst phenomenon. For example, *Mandelbrot and Van Ness* [1968] introduced a fractional Brownian noise model, which can produce stationary streamflows with a specified Hurst coefficient. After these findings by *Hurst* [1951] and *Mandelbrot and Van Ness* [1968], there emerged a considerable debate in the water resource literature regarding the implications of the Hurst phenomenon. *Klemes* [1974] argued that our understanding of long-term persistence, and even stochastic hydrology in general, is (was) at about the same stage as the Ptolemaic planetary model prior to the introduction of Kepler's laws of planetary motion. On the one hand, *Wallis and O'Connell* [1973, p. 363] suggest that "to emphatically state that hydrologic records do not exhibit long-term persistence demands either a very naïve understanding of statistics or a monumentally large database." On the other hand, *Klemes et al*. [1981, p. 737] argue that in spite of all our knowledge gained from the 1970s literature on the Hurst phenomenon, "the replacement of short-memory models with long-memory models in reservoir analyses cannot be objectively justified." It seems fitting that we should at least attempt to resurrect this debate by using the large database used in the previous sections of this study [*Slack et al*., 1993] to explore the behavior of the Hurst coefficient of annual streamflow traces in the continental United States.

4.1. Estimation of the Hurst Coefficient

Montanari et al. [1997] review numerous methods for estimation of the Hurst coefficient H and numerous sources of literature relating to the detection of long memory. *Hurst* [1951] suggested simply plotting the logarithm of the computed values of the rescaled range R_n/s_x versus the logarithm of the sample size n, and computing the Hurst coefficient as the slope of the resulting empirical relation. *Mandelbrot and Wallis* [1995] describe the deficiencies of this approach. *Mandelbrot and Wallis* [1969, 1995] suggested repeating this graphical approach for many subsamples of size $k \leq n$ from the same record of length, leading to a pox diagram. A pox diagram is simply a log-log plot of the rescaled range R_n/s_x versus subsample length. An estimate of H is then derived as the slope of the pox diagram. *Mandelbrot and Wallis* [1995] provide many examples of pox diagrams along with a discussion of their properties. In this study, pox diagrams for each of the 1544 sites were constructed, and the resulting estimates of H obtained via regression for each site are summarized in Figure 7 as a function of the record length associated with each record available in the *Slack et al*. [1993] database. The locally weighted scatterplot smooth (LOWESS) shown in Figure 7 illustrates the well-known effect of record length on estimates of the Hurst coefficient. Small sample upward bias associated with estimates of H illustrated in Figure 7 is a well-documented phenomenon [*Wallis and Matalas*, 1970]. The upward bias in estimators of H increases as n decreases and as ρ_1 increases and is generally larger for normally distributed variables than for nonnormal variates [*Wallis and Matalas*, 1970].

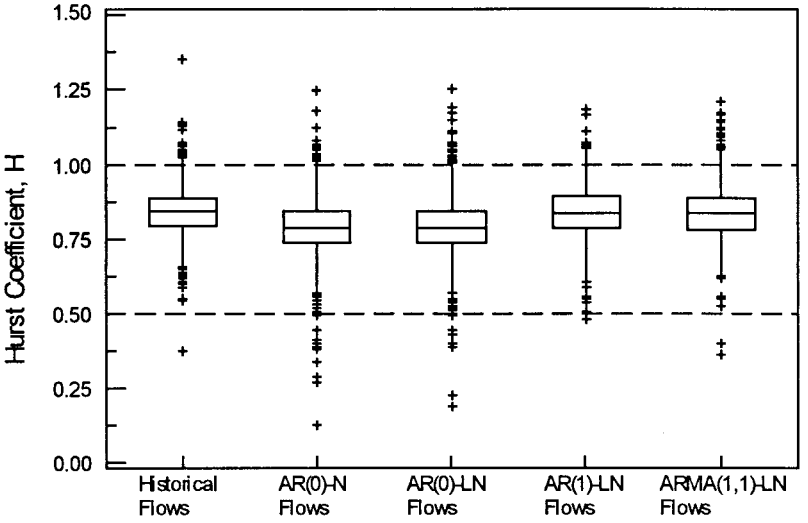

Figure 8. Box plots illustrating the distribution of Hurst coefficients estimated from historical records and from synthetic flow sequences generated by four different models.

4.2. The Implications of the Hurst Phenomenon and Estimates of *H*

There is still considerable controversy and debate regarding the implications of the Hurst phenomenon. Although many investigators have computed Hurst coefficients for actual streamflow and other geophysical records (see *Mandelbrot and Wallis* [1995] and *Montanari et al*. [1997] for summaries and citations), there is considerable literature that questions the meaning of computed values of H. For example, *Potter* [1979] documents that nonhomogeneity of hydrologic records can lead to inflation of estimated autocorrelation functions and Hurst coefficients. Similarly, *Klemes* [1974] shows that obtaining values of $H > 0.5$ is not necessarily an indication that the process has infinite memory. Klemes shows that the Hurst phenomenon can result from nonstationarity or from storage processes which arise in many natural systems.

Using a set of 1009 streamflow records developed by *Wallis et al*. [1991], *Lettenmaier et al*. [1994] document that strong trends exist for a large number of these basins, many more than would be expected due to chance. *Lins* [1985a] also documents that annual streamflow in the United States is probably not stationary. This study uses a different database of streamflow [*Slack et al*., 1993] with records over the period 1874–1988. We expect that the estimates of H in Figure 7 are due, in part, to nonstationarity and even nonhomogeneity of the flow records. The extensive literature on the properties of the Hurst coefficient and the Hurst phenomenon demonstrates that the estimates of $H > 0.5$ in Figure 7 do not constitute evidence of long-term persistence. Without definitive proof that flow records are stationary and homogeneous, it is impossible to use H to infer properties regarding the long-term persistence structure of time series. Instead, we use H as a goodness-of-fit measure in an attempt to select a stochastic structure which resembles the historical flow series.

4.3. Experimental Design: Using *H* to Select a Stochastic Structure for Annual Streamflow Records

In this section we describe experiments which use H to determine which stochastic process best approximates sequences of annual streamflow in the United States. These experiments consider both type I and type II errors. The first set of type I experiments compares estimated Hurst coefficients obtained from the 1544 historical flow sequences with estimated Hurst coefficients obtained from 1544 synthetic flow sequences derived from each of four synthetic streamflow models. The four synthetic models include the following ARMA models: AR(0)-N, AR(0)-LN, AR(1)-LN, and ARMA(1,1)-LN, where N denotes normal distribution and LN denotes the lognormal distribution. Equation (1) is an AR(1)-LN model. We include an AR(0)-N to evaluate the influence of the marginal distribution on estimation of H. *Loucks et al*. [1981] provide an introduction to ARMA models. We employ the approach described by *Lettenmaier and Burges* [1977] to generate flow sequences from an ARMA(1,1)-LN model.

Our first set of experiments are type I experiments because our null hypothesis, in each case, is that the historical flows arise from that particular synthetic flow generator. These type I experiments are analogous to examining type I errors in hypothesis testing. Unfortunately, type II errors are also possible. Our second experiments evaluate type II errors; that is, if the historical flows really do arise from an AR(1)-LN model, as is hypothesized in (1), would we find an ARMA(1,1)-LN model acceptable (type II error) and vice versa.

4.3.1. Type I experimental results. Initially, 1544 synthetic flow sequences were generated from AR(0)-N, AR(0)-LN, AR(1)-LN, and ARMA(1,1)-LN models, with model parameters defined by the 18 regional mean values of C_v, ρ_1, and ρ_2. Unbiased estimators of C_v and ρ_1 are obtained using $\hat{C}_{v3}$ in (18) and r_1^* in (10), respectively. An unbiased estimator of ρ_2 is derived in the appendix. The record length of each synthetic flow sequence is equal to one of the historical records in that region, so that across each region, the distribution of the record lengths of the synthetic flow sequences is identical to that of the historical records. Figure 8 compares box plots of estimates of H obtained from the historical flow sequences and the four synthetic flow generators. We conclude from a comparison of the box plots in Figure 8 that in terms of the distribution of estimated Hurst coefficients, the historical flows could not be distinguished from the synthetic AR(1)-LN and

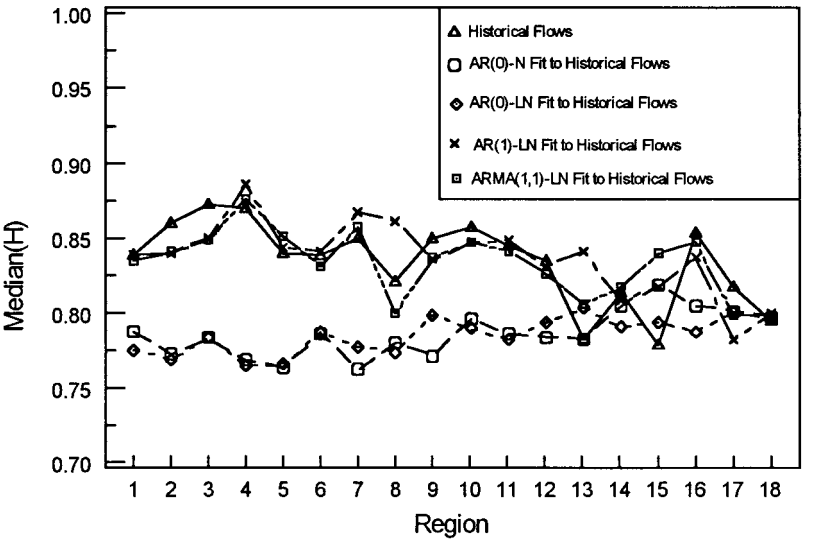

Figure 9. The median Hurst coefficient estimated from historic flow sequences and from four stochastic streamflow models fit to each region.

the ARMA(1,1)-LN flows. Figure 8 also illustrates that an AR(0) model cannot reproduce the distribution of the historical estimates of the Hurst coefficient. Figure 8 further illustrates that the marginal distribution of the flows does not seem to have a significant impact on the distribution of estimated Hurst coefficients, because the results for the AR(0)-N and AR(0)-LN models are identical.

Figure 9 illustrates the median value of H estimated for each of the 18 regions of the United States shown in Figure 1. Figure 9 documents, by region, that the AR(1)-LN and ARMA(1,1)-LN models reproduce the historical median values of H in most regions. We conclude from these initial graphical comparisons that the historical streamflow sequences are indistinguishable, in terms of H, from the AR(1)-LN and ARMA(1,1)-LN synthetic flow sequences.

4.3.2. Type II experimental results. It remains to determine whether the Hurst coefficient can be used to distinguish between AR(1)-LN and ARMA(1,1)-LN processes. To answer this question, the same experiment as above is repeated; however, now pseudohistorical flow sequences are created using the synthetic flow models. The question we address is whether we would accept flow sequences as being AR(1)-LN when they are actually generated by an ARMA(1,1)-LN process and vice versa.

Figures 10 and 11 are both identical to Figure 9 except that the historical flow sequences are replaced by AR(1)-LN and ARMA(1,1)-LN synthetic flow traces, respectively. Figure 10 documents that when pseudohistoric flow sequences are AR(1)-LN, both the AR(1)-LN and the ARMA(1,1)-LN models, both reproduce the median Hurst coefficients estimated from those AR(1)-LN flow sequences. This is expected because an AR(1) model is simply a special case of an ARMA(1,1) model. What are interesting, however, are the results in Figure 11 which document that even when

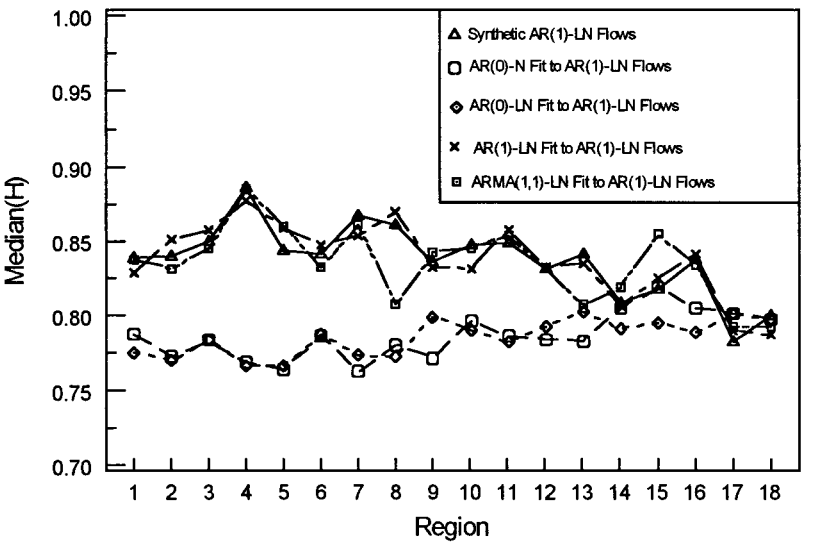

Figure 10. The median Hurst coefficient estimated from pseudohistoric AR(1)-LN flow sequences and from four stochastic streamflow models fit to those sequences.

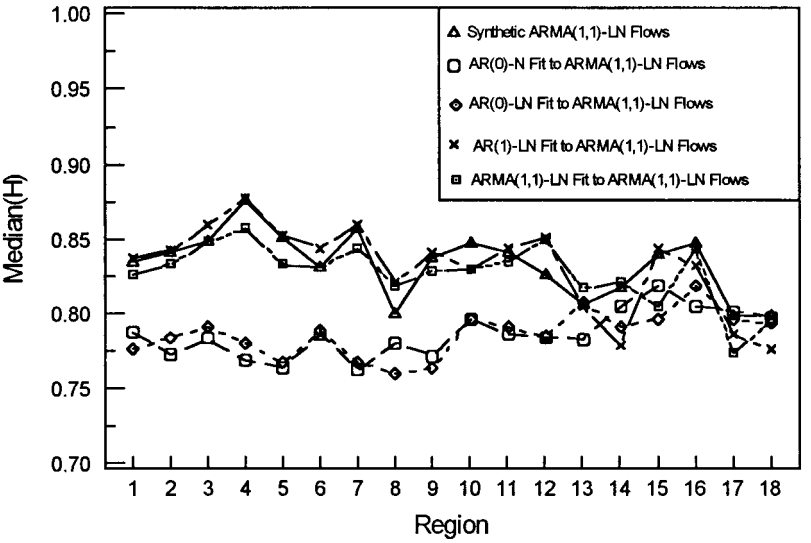

Figure 11. The median Hurst coefficient estimated from pseudohistoric ARMA(1,1)-LN flow sequences and from four stochastic streamflow models fit to those sequences.

pseudohistoric flows are generated by an ARMA(1,1) model structure, we are unable to discern whether those flows are either AR(1) or ARMA(1,1); in fact, they could be either. We conclude from these experiments that even if historical flows were ARMA(1,1), we would be unable to distinguish their stochastic structure from an AR(1) model, so that, in practice, either model is acceptable. If either model is acceptable, though, it is well-established practice to choose the more parsimonious one. We conclude that an AR(1)-LN model provides an acceptable model of the structure of observed historical flow sequences because it is able to reproduce the observed distribution of Hurst coefficients. *Tsai* [1997] provides more definitive comparisons using quantitative hypothesis tests to prove that the results reported in Figures 8–11 are truly representative and not due to chance.

5. Conclusions

This study has sought to improve our understanding of the spatial and temporal structure of historical annual streamflow in the continental United States. Using a large database of annual streamflows developed by *Slack et al.* [1993], we document that the observed variability in estimates of the lag-one serial correlation coefficient in each of 18 major water resource regions could easily be due to sampling alone. This result implies that large differences in estimates of ρ do not imply large differences in the physical processes that give rise to persistence, as is often assumed. Instead, the large differences in at-site estimates of ρ are probably due to sampling alone. These results emphasize the importance of exploiting regional information regarding persistence rather than trusting at-site data alone. However, the 18 broad geographic regions are not homogeneous in terms of the coefficient of variation of annual streamflow. To assure homogeneity in terms of the year-to-year variability of streamflow, we document that the nation must be broken into over 200 regions. This result is quite useful for studies which seek to perform national or regional hydrologic assessments. Regional assessments based on only 18 water resource regions may not be useful because most of those regions are heterogeneous in terms of hydrologic variability. Instead, at least 106 assessment subregions are required, as in the second national assessment [*U.S. Water Resources Council*, 1975], or better yet, the 222 planning subregions introduced by the U.S. Water Resources Council are required to assure hydrologic homogeneity.

In addition to these empirical results, this study also summarized the sampling properties of estimates of C_v and ρ. An unbiased estimator of C_v for an autoregressive lognormal (AR(1)-LN) process was derived which depends on sample size and lag-one correlation ρ. This estimator was used to obtain unbiased estimates of regional mean C_v. As is often the case, the unbiased estimator $\hat{C}_{v3}$ introduced here had slightly higher variance than either of the biased alternatives, yet still it is recommended for applications such as this one, where one's interest is in obtaining an unbiased regional estimate of C_v.

Finally, experiments were performed to evaluate whether various simple stochastic streamflow models could generate flow traces with Hurst coefficients which resemble historical flow sequences. Our experiments document that both a lognormal autoregressive (AR(1)-LN) and a lognormal autoregressive moving average (ARMA(1,1)-LN) model can reproduce the distribution of estimates of Hurst (H) coefficients across the entire continental United States. We also document that if the historical streamflows really were generated from an ARMA(1,1)-LN process, estimates of H appeared identical to those obtained from an AR(1)-LN model and vice versa. These results indicate that even if natural flow records do exhibit long-term persistence (an ARMA(1,1) model can exhibit infinite memory), we are unable to distinguish that memory from AR(1)-LN flow sequences which cannot exhibit long-term persistence. These results only confirm our inability to discern the complex long-term persistence structure of natural flow records due to the fact that only short records are available. Assessment of the long-term persistence structure of actual flow records is further confounded by nonstationarity and nonhomogeneity of those flow records.

Appendix

A1. Derivation of Var[$\hat{C}_{v3}$] for an AR(1)-LN Process

Assuming streamflow x is from the AR(1)-LN process given in (1), we derive the Var[$\hat{C}_{v3}$], where $\hat{C}_{v3} = \sqrt{\exp(v_y^{*2}) - 1}$. Approximating $\hat{C}_{v3}$ with a first-order Taylor series about the true variance σ_y^2 leads to

$$\hat{C}_{v3} \cong \sqrt{\exp(\sigma_y^2) - 1} + \frac{d\sqrt{\exp(\sigma_y^2) - 1}}{d\sigma_y^2}[v_y^{*2} - \sigma_y^2] \tag{A1}$$

Now using the variance operator leads to

$$\mathrm{Var}[\hat{C}_{v3}] \cong \left(\frac{d\sqrt{\exp(\sigma_y^2) - 1}}{d\sigma_y^2}\right)^2 \mathrm{Var}[v_y^{*2}] \tag{A2}$$

which can be simplified to

$$\mathrm{Var}[\hat{C}_3] \cong \left[\frac{1 + C_v^2}{2\theta C_v}\right]^2 \mathrm{Var}[v_y^2] \tag{A3}$$

where

$$\theta = \left[1 + \frac{1}{n} + \frac{2[\rho_{y,1}(n+1) - n]}{n^2(1 - \rho_{y,1})^2}\right].$$

Since the logarithms of streamflow y follow an AR(1) normal model, *Loucks et al*. [1981, equation 3C.9] document that

$$\mathrm{Var}[v_y^2] \cong \frac{2\sigma_y^4}{n}\left[\frac{1 + \rho_{y1}^2}{1 - \rho_{y1}^2}\right] \tag{A4}$$

Combining (A3) and (A4) and using the fact that $\sigma_y^2 = \ln(1 + C_v^2)$ leads to the result in (19). Figure A1 illustrates that the theoretical approximation in (19) provides an excellent approximation to experimental estimates of the standard error of the estimator $\hat{C}_{v3}$ based on Monte Carlo experiments. Here standard error (SE) is defined as $\mathrm{SE}(\hat{C}_{v3}) = \sqrt{\mathrm{Var}[\hat{C}_{v3}]}$.

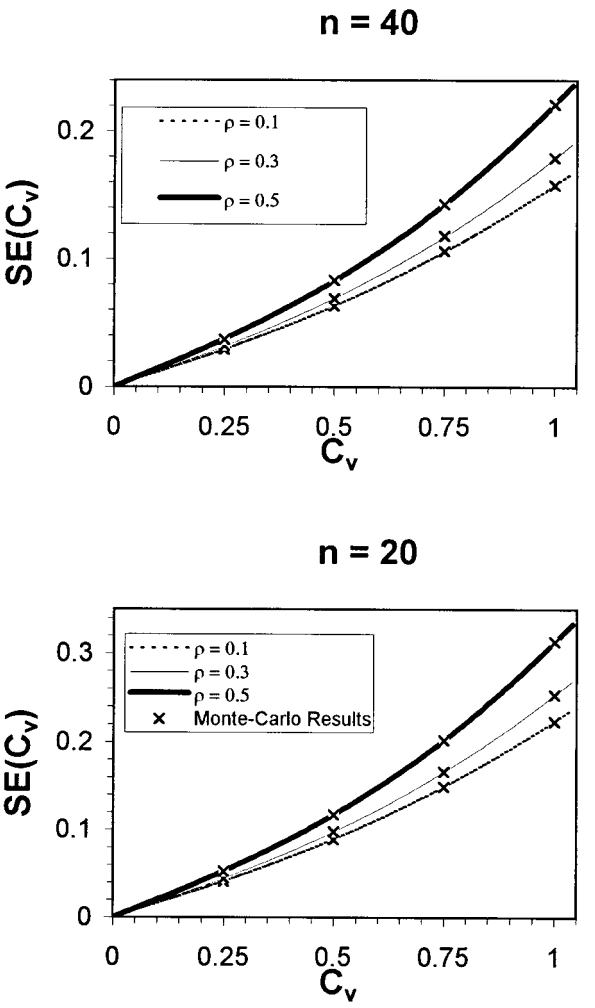

Figure A1. Comparison of theoretical and experimental estimates of the standard error of $\hat{C}_{v3}$.

A2. Derivation of an Unbiased Estimator of ρ_2

The estimator r_k in (7) is a biased estimator of ρ_k. *Kendall* [1954] derived the expectation of r_k for an AR(1) process:

$$E[r_k] = \rho_1^k - \frac{1}{N - k}\left\{\frac{1 + \rho_1}{1 - \rho_1}(1 - \rho_1^k) + 2k\rho_1^k\right\} \tag{A5}$$

We used the same approach used by *Wallis and O'Connell* [1972] to derive the following unbiased estimator of ρ_2, which we term r_2^*. Setting $k = 2$ in (A5), replacing $E[r_2]$ in (A5) with r_2 obtained from (7), and then solving (A5) for $r_1^{**} = \rho_1$ leads to

$$r_1^{**} = \frac{2 \pm \sqrt{4 + 4(n-7)[1 + (n-2)r_2]}}{2(n-7)} \tag{A6}$$

Now an unbiased estimator of ρ_2 is obtained by using the fact that the correlogram for an AR(1) process is $\rho_k = \rho_1^k$; hence

$$r_2^* = (r_1^{**})^2 \tag{A7}$$

Monte Carlo experiments documented that for short records such as those used in this study, the smaller value of r_1^* in (A6) is preferred. The regional mean values of r_1^* and r_2^* are reported in Table 1. The regional mean values of r_1^* in Table 1 differ from the regional mean values of $r_{y,1}^*$ in Table 1 because they represent the lag correlations in real and log space, respectively.

Acknowledgments. Although the research described in this paper has been funded in part by the United States Environmental Protection Agency through grant R 824992-01-0 to Tufts University, it has not been subjected to the Agency's required peer and policy review and therefore does not necessarily reflect the views of the Agency, and no endorsement should be inferred. The authors are indebted to Gary Tasker, John Dracup, and an anonymous reviewer for their constructive comments on an earlier version of this manuscript.

References

Bartlein, P. J., Streamflow anomaly patterns in the U.S.A. and southern Canada—1951–1970, *J. Hydrol*., *57*, 49–63, 1982.

Bobée, B., and P. F. Rasmussen, Recent advances in flood frequency analysis, *U.S. Natl. Rep. Int. Union Geod. Geophys. 1991–1994*, *Rev. Geophys*., *33*, 1111–1116, 1995.

Fill, H. D., and J. R. Stedinger, Homogeneity tests based upon Gumbel distribution and a critical appraisal of Dalrymple's test, *J. Hydrol*., *166*, 81–105, 1995.

Guetter, A. K., and K. P. Georgakakos, River outflow of the conterminous United States, 1939–1988, *Bull. Am. Meteorol. Soc*., *74*(10), 1873–1891, 1993.

Hosking, J. R. M., and J. R. Wallis, Regional Frequency Analysis, Cambridge Univ. Press, New York, 1997.

Hurst, H. E., Long-term storage capacity of reservoirs, *Trans. Am. Soc. Civ. Eng.*, *116*, 770–808, 1951.

Jennings, M. E., W. O. Thomas Jr., and H. C. Riggs, Nationwide summary of U.S. Geological Survey regional regression equations for estimating magnitude and frequency of floods for ungaged sites, 1993, *U.S. Geol. Surv. Water Resour. Invest. Rep. 94-4002*, 196 pp., 1994.

Kendall, M. G., Note on the bias in the estimation of autocorrelation, *Biometrika*, *42*, 403–404, 1954.

Klemes, V., The Hurst phenomenon: A puzzle?, *Water Resour. Res.*, *10*(4), 675–688, 1974.

Kirby, W., Algebraic boundedness of sample statistics, *Water Resour. Res.*, *10*(2), 220–222, 1974.

Klemes, V., R. Srikanthan, and T. A. McMahon, Long-memory flow models in reservoir analysis: What is their practical value?, *Water Resour. Res.*, *17*, 737–751, 1981.

Leipnik, R. B., Distribution of the serial correlation coefficient in a circulatory correlated universe, *Ann. Math. Stat.*, *18*, 80–87, 1947.

Lettenmaier, D. P., and S. J. Burges, An operational approach to preserving skew in hydrologic models of long-term persistence, *Water Resour. Res.*, *13*(2), 281–290, 1977.

Lettenmaier, D. P., E. F. Wood, and J. R. Wallis, Hydro-climatological trends in the continental United States, 1948–1988, *J. Clim.*, *7*, 586–607, 1994.

Lins, H. F., Streamflow variability in the United States, *J. Clim. Appl. Meteorol.*, *24*, 463–471, 1985a.

Lins, H. F., Interannual streamflow variability in the United States based on principal components, *Water Resour. Res.*, *21*(5), 691–701, 1985b.

Lins, H. F., Regional streamflow regimes and hydroclimatology of the United States, *Water Resour. Res.*, *33*(7), 1655–1667, 1997.

Loucks, D. P., J. R. Stedinger, and D. A. Haith, *Water Resource Systems Planning and Analysis*, Prentice-Hall, Englewood Cliffs, N. J., 1981.

Mandelbrot, B. B., and J. W. Van Ness, Fractional Brownian motions, fractional Brownian noise and applications, *SIAM Rev.*, *10*(4), 422–437, 1968.

Mandelbrot, B. B., and J. R. Wallis, Robustness of the rescaled range R/S in the measurement of noncyclic long run statistical dependence, *Water Resour. Res.*, *5*(5), 967–988, 1969.

Mandelbrot, B. B., and J. R. Wallis, Some long-run properties of geophysical records, in *Fractals in the Earth Sciences*, edited by C. C. Barton and P. R. LaPointe, chap. 2, pp. 41–64, Plenum, New York, 1995.

Markovic, R. D., Probability of best fit to distributions of annual precipitation and runoff, *Hydrol. Pap. 8*, Colo. State Univ., Fort Collins, 1965.

Montanari, A., R. Rosso, and M. S. Taqqu, Fractionally differenced ARIMA models applied to hydrologic time series: Identification, estimation, and simulation, *Water Resour. Res.*, *33*(5), 1035–1044, 1997.

Potter, K. W., Annual precipitation in the northeast United States: Long memory, short memory, or no memory?, *Water Resour. Res.*, *15*(2), 340–346, 1979.

Slack, J. R., A. M. Lumb, and J. M. Landwehr, Hydro-Climatic Data Network (HCDN): Streamflow data set, 1874–1988, [CD-ROM], *U.S. Geol. Surv.*, *Water Resour. Invest. Rep. 93-4076*, 1993.

Stedinger, J. R., Estimating correlations in multivariate streamflow models, *Water Resour. Res.*, *17*(1), 200–208, 1981.

Stuart, A., and J. K. Ord, Standard errors, in *Kendall's Advanced Theory of Statistics*, 5th ed., vol. 1, chap. 10, pp. 243–262, Oxford Univ. Press, New York, 1987.

Tasker, G. D., Approximate sampling distribution of the serial correlation coefficient for small samples, *Water Resour. Res.*, *19*(2), 579–582, 1983.

Tsai, Y.-S., The long term persistence of annual streamflow in the United States, M.S. thesis, Tufts Univ., Medford, Mass., 1997.

U.S. Water Resources Council, *The Nation's Water Resources 1975–2000, Second National Assessment by the U.S. Water Resources Council*, vol. 3, *Analytical Data*, *Appendix V, Streamflow Conditions*, 276 pp., U.S. Gov. Print. Off., Washington, D. C., 1975.

Vogel, R. M., and N. M. Fennessey, *L* moment diagrams should replace product moment diagrams, *Water Resour. Res.*, *29*(6), 1745–1752, 1993.

Vogel, R. M., and I. Wilson, Probability distribution of annual maximum, mean and minimum streamflows in the United States, *J. Hydrol. Eng.*, *1*(2), 69–76, 1996.

Wallis, J. R., and N. C. Matalas, Small sample properties of H and K: Estimators of the Hurst coefficient *h*, *Water Resour. Res.*, *6*(6), 1583–1594, 1970.

Wallis, J. R., and P. E. O'Connell, Small sample estimation of ρ_1, *Water Resour. Res.*, *8*(3), 707–712, 1972.

Wallis, J. R., and P. E. O'Connell, Firm reservoir yield—How reliable are historic records?, *Hydrol. Sci. Bull.*, *18*(3), 347–365, 1973.

Wallis, J. R., D. P. Lettenmaier, and E. F. Wood, A daily hydroclimatological data set for the continental United States, *Water Resour. Res.*, *27*(7), 1657–1663, 1991.

J. F. Limbrunner, Y. Tsai, and R. M. Vogel, Department of Civil and Environmental Engineering, Tufts University, 113 Anderson Hall, Medford, MA 02115. (e-mail: rvogel@tufts.edu)

(Received October 6, 1997; revised July 27, 1998;
accepted July 28, 1998.)